SHOCKS AND GOVERNMENT BELIEFS: THE RISE AND FALL OF AMERICAN INFLATION

THOMAS SARGENT, NOAH WILLIAMS, AND TAO ZHA

ABSTRACT. We use a Bayesian Markov Chain Monte Carlo algorithm jointly to estimate the parameters of a 'true' data generating mechanism and those of a sequence of approximating models that a monetary authority uses to guide its decisions. Gaps between a true expectational Phillips curve and the monetary authority's approximating non-expectational Phillips curve models unleash inflation that a monetary authority that knows the true model would avoid. A sequence of dynamic programming problems implies that the monetary authority's inflation target evolves as its estimated Phillips curve moves. Our estimates attribute the rise and fall of post WWII inflation in the US to an intricate interaction between the monetary authority's beliefs and economic shocks. Shocks in the 1970s made the monetary authority perceive a tradeoff between inflation and unemployment that ignited big inflation. The monetary authority's beliefs about the Phillips curve changed in ways that account for Volcker's conquest of US inflation.

Date: September 6, 2004; Revised April 24, 2005.

Key words and phrases. History dependence; updating beliefs; policy evaluation; self-confirming equilibrium; Nash inflation; Ramsey outcome.

We thank seminar participants at the NBER Summer Institute, Eric Leeper, James Nason, Ricardo Reis, William Roberds, Frank Schorfheide, Christopher Sims, and Lars Svensson for helpful discussions and Jordi Mondria for excellent research assistance. We also thank two anonymous referees for detailed and helpful comments and suggestions. The views expressed herein do not necessarily reflect those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Sargent and Williams thank the National Science Foundation for (separate) research support.

I. INTRODUCTION

Today, many statesmen and macroeconomists believe that inflation can largely be determined by a government monetary authority. Then why did the Federal Reserve Board preside over high US inflation during the late 1960s and the 1970s? And why, under Paul Volcker, did it quickly arrest inflation during the early 1980s? This paper answers these questions by estimating a model that allows discrepancies between a true data generating mechanism and a monetary's authority approximating model. Our model features a process that makes a sequence of economic shocks induce the monetary authority to alter its model of inflation-unemployment dynamics, the Phillips curve. At each date t, the monetary authority updates its beliefs about the Phillips curve and then computes a first-period action recommended by a "Phelps problem", a discounted dynamic programming problem that minimizes the expected value of a discounted quadratic loss function of inflation and unemployment.[1] The monetary authority pursues the same objectives at each date and uses the same structural model. But its estimates of that model change.[2] This model of the systematic part of inflation puts the monetary authority's beliefs about the Phillips curve front and center.[3]

We assume that the monetary authority's model of the Phillips curve deviates in two ways from the true data generating model, a version of Robert E. Lucas Jr.'s (1973) aggregate supply function used by Kydland and Prescott (1977) and many others. The first deviation is that the monetary authority omits the public's rational expectation of inflation from its Phillips curve. By itself, this omission need not prevent the outcomes of our model from coinciding with those predicted by Kydland and Prescott, nor need it imply that the government's model is wrong in a way that could be detected even from an infinite sample.

---

[1] Sargent (1999) called it a Phelps problem.

[2] There is some debate about whether policy objectives or the structural models used by policymakers have evolved over time. However introducing such an evolution of understanding into formal models is difficult without arbitrarily imposing exogenous changes. We need no such exogenous shifts.

[3] As does Kydland and Prescott's (1977) model of time-consistent suboptimal inflation.

The reason is that, depending on the history of outcomes, the constant term and lagged rates of inflation and unemployment can stand in perfectly for the expected rate of inflation that the government has omitted from its Phillips curve.[4] If the monetary authority were to believe that the coefficients of its Phillips curve were constant over time, then its estimates would converge to ones that support a self-confirming equilibrium (SCE). After convergence, its estimated Phillips curve would correctly describe occurrences along the SCE path for inflation and unemployment. Such an after-convergence version of our model has little hope of explaining the rise and fall of US inflation: that model would have inflation fluctuating randomly around a constant SCE level that coincides with Kydland and Prescott's time consistent suboptimal (i.e., excessive) level.[5]

This outcome motivates our second subtle deviation from a rational expectations equilibrium. Instead of thinking that the regression coefficients in its Phillips curve are time invariant (which they indeed *are* in an SCE), our monetary authority believes that they form a vector random walk with innovation covariance matrix V. Given that model, the monetary authority updates its beliefs using Bayes' rule. The covariance matrix V and the initial condition for the regression parameters in the monetary authority's Phillips curve become hyperparameters of a model that shapes evolution of the monetary authority's beliefs.[6] After calibrating the initial condition and imposing that the systematic part of inflation is determined by the time t solution of the Phelps problem, we estimate V along with parameters of the *true* expectational Phillips curve that, unbeknownst to the monetary authority, actually governs inflation-unemployment dynamics. We use a Bayesian Markov

[4]See Kydland and Prescott (1977) for a heuristic argument and Sargent (1999) for a demonstration that the outcome in a self-confirming equilibrium is identical with Kydland and Prescott's time-consistent outcome.

[5]Parkin (1993) and Ireland (1999) advocate the hypothesis that the post WWII US inflation data can be accounted for by well understood medium term movements in the natural rate of unemployment, stable government preferences, and steady adherence to the time-consistent suboptimal equilibrium of Kydland and Prescott (1977).

[6]As is true in a rational expectations model, the monetary authority's beliefs are outcomes, not free parameters.

Chain Monte Carlo (MCMC) algorithm to estimate statistics that describe the posterior distribution of these parameters of our model. We obtain a much better explanation of the monetary authority's inflation choices than earlier efforts to estimate similar models had achieved.

We find that our model fits the data better than a benchmark time series model. We use several criteria to compare the empirical performance of our theoretical model with those of some atheoretical Bayesian vector autoregressions (BVARs). Our model has better forecasting performance than BVARs over one-month, two-year, and four-year horizons. Formal model selection criteria, such as the Schwarz criterion and Bayes factors, strongly favor our model. Equally important, our model outperforms the BVARs in predicting several key turning points in the inflation time series. Finally, while the fit of our model is competitive with statistical models, our results yield important insights that help to *understand* the US inflation experience, something a purely statistical model cannot. One essential feature accounting for our model's success in fitting the data is how our estimation procedure exploits the cross-equation restrictions that the government's Phelps problem imposes on the sequence of government beliefs about the empirical Phillips curve. These restrictions are very informative for estimating the key government belief parameters in V.

With particular a priori settings of the parameter innovation covariance matrix V, Sims (1988), Chung (1990), Sargent (1999), and Cho, Williams, and Sargent (2002) all studied versions of our model.[7] When Chung and Sargent estimated their a-priori-fixed-V versions of our model, they obtained discouraging results. They did not come close to explaining the rise and fall of US inflation in terms of a process of the monetary authority's learning about its Phillips curve.[8]

---

[7]Sargent and Williams (2005) is an extended theoretical study of a version of our model that focuses on the impact of different settings of V on rates of convergence to, escapes from, and cycles around an SCE.

[8]Previous failures to match the data with a model like ours seem to be widely recognized and helped to promote a literature that makes the "stickiness" (or persistence) of inflation exogenous.

This paper estimates settings for V that attain substantial improvements in the model's ability to rationalize the choices made by the US monetary authority. The MCMC algorithm finds values for V that allow the model to reverse engineer a sequence of government beliefs about the Phillips curve that, through the intermediation of the Phelps problem, capture both the acceleration of US inflation in the 1970s and its rapid decline in the early 1980s. Our MCMC method estimates a V that accommodates an avenue by which economic shocks impinge on the monetary authority's beliefs, via its use of Bayes' rule, and its decisions, via successive solutions of its Phelps problem. The monetary authority's views about parameter drift and its application of Bayes' rule add a source of history dependence to its procession of decisions that is absent in either the SCE or the Markov perfect equilibrium of Kydland and Prescott's model. The resulting interactions of shocks and monetary beliefs forms the basis for our explanation of the rise and fall of US inflation.

The rest of the paper is organized as follows. Section II relates our findings to other work. In Section III, we lay out the model and discuss theoretical characterizations of it. Section IV develops an econometric methodology for estimation, and Section V reports the estimated results. In Section VI, we present further empirical results, stress the importance of cross-equation restrictions via the Phelps problem, examine the forecasting performance of the model, conduct some counterfactual exercises, and explore some important implications. Section VII discusses the estimated model's long-run properties. Section VIII concludes. Four appendices describe the data and provide technical details about our prior distribution and the posterior sampling scheme.

II. Relation to Recent Literature

Cogley and Sargent's (2005a) explanation of US post WWII inflation also features the interaction of a government learning process and a sequence of Phelps problems. Cogley and Sargent's government applies Bayes' rule recurrently to estimate three Phillips curve models, only one of which is a rational expectations version of a natural rate model. Cogley and Sargent focus on the role of model uncertainty in policy making and take no stand on the true data generating mechanism. Their story is about how an almost discredited model that has even a very small posterior probability will nevertheless be very influential if it leads to very bad outcomes under the policies that would have been recommended if its posterior probability were *exactly* zero.[9]

Primiceri (2003) also develops a learning model to explain the rise and fall of US inflation. He estimates his model on US data and finds that its fit is comparable to an atheoretical VAR as a description of the data. Like us, he emphasizes that inflation remained high in the 1970s due to the government's perception that disinflation was too costly. Unlike us, a key component of his story is that the monetary authority's mismeasurement of the natural rate of unemployment caused policy to be looser than policymakers intended. Primiceri's main focus is on a backward-looking Keynesian model that has no explicit role for private sector expectations that respond to the government's decision rule, unlike our true model, which has private sector expectations responding in the best way to government decisions.[10] In the SCE of Primiceri's model, the government learns that the coefficients on inflation in the Phillips curve sum to one. This restriction is correct in his specification, but need not hold under rational expectations.[11] Primiceri finds that the model's transient learning dynamics are able to reproduce low frequency features of post WWII US inflation-unemployment dynamics.

Our paper differs substantially from Primiceri's because our true data generating mechanism (DGM) is a rational expectations natural rate model. Policymakers' misspecified

[9]The monetary authority in our model updates only one model, an outmoded one at that, in light of the rational expectations revolution. From their readings of minutes of the FOMC, Christina Romer and David Romer (2002) infer that the Fed's learning process was confined to a primitive Phillips curve specification like the one we impute to the monetary authority. Their story about the evolution of Fed beliefs assigns no influence to rational expectations ideas.

[10]Primiceri also considers a New Keynesian rational expectations model, but it fits substantially worse than his backward-looking specification.

[11]See Sargent (1999) for a discussion.

model can eventually converge to attain a self-confirming equilibrium in which the inflation-unemployment dynamics generated by the true DGM agree with those expected by the government along the equilibrium outcome process. Unlike Primiceri, we view the rational expectations natural rate theory and the associated SCE as a useful starting point. We build on Sims (1988), Chung (1990), Sargent (1999), and Cho, Williams, and Sargent (2002), as generalized by Sargent and Williams (2005). These studies a priori adopted parameter specifications that opened a substantial gap between a Ramsey inflation outcome (the one that would be chosen by a government that knew the correct DGM) and the Nash inflation outcome that emerges from the SCE. The latter three contributions discovered *mean dynamics* that on average push outcomes toward the Nash inflation level and *escape dynamics* that recurrently push them toward the Ramsey outcome.

The present paper estimates key parameters that control the mean dynamics and the escape dynamics. Our empirical estimates teach us to deemphasize the empirical relevance of *both* the mean dynamics and the escape dynamics and instead to focus on the short-term impacts of shocks on government beliefs. In addition, our estimate of a small gap between the Nash and Ramsey inflation levels supports Blinder's (1998) skepticism about whether that gap is quantitatively important for the monetary authority's decision problem.

III. The Model

We extend the model of Sargent and Williams (2005). There is a Lucas natural-rate version of the Phillips curve and a true inflation process:

$$u_t - u^* = \theta_0(\pi_t - E_{t-1}\pi_t) + \theta_1(\pi_{t-1} - E_{t-2}\pi_{t-1}) + \tau_1(u_{t-1} - u^*) + \sigma_1 w_{1t}, \qquad (1)$$

$$\pi_t = x_{t-1} + \sigma_2 w_{2t}, \qquad (2)$$

where u_t is the unemployment rate, π_t is inflation, x_t is the part of inflation controllable by the government given the information up to time t, and w_{1t} and w_{2t} are i.i.d. uncorrelated standard normal random variables. Equation (1) is an expectations-augmented Phillips curve in which systematic monetary policy has neither short-run nor long-run effects on unemployment.[12] Equation (1) embodies a stronger form of 'policy irrelevance' than do many of today's New Keynesian Phillips curves. In this paper, we ignore the nonneutralities present in those models and aim to reverse engineer a set of government beliefs that can explain the low frequency swings in U.S. data while insisting that the true DGM have the strong policy irrelevance of the Lucas supply function. Section V shows that our reverse-engineering succeeds quantitatively in tracking the post-WWII inflation data.

Equation (2) states that the government determines inflation up to a random shock. The public has rational expectations, so that $E_{t-1}\pi_t = x_{t-1}$. The government dislikes inflation and unemployment. The policy decision x_{t-1} solves the "Phelps problem":

$$\min_{x_{t-1}} \hat{E} \sum_{t=1}^{\infty} \delta^t ((\pi_t - \pi^*)^2 + \lambda (u_t - u^{**})^2) \tag{3}$$

subject to (2) and

$$u_t = \hat{\alpha}'_{t|t-1}\Phi_t + \sigma w_t, \tag{4}$$

where π^* and u^{**} are the targeted levels of inflation and unemployment, both $\hat{\alpha}_{t|t-1}$ and Φ_t are $r \times 1$ vectors, w_t is an i.i.d. standard normal random variable, and where (4) is the monetary authority's model of inflation-unemployment dynamics. The vector Φ_t of regressors consists of lags of unemployment and inflation. By comparing (4) with the true DGM (1), we see that the government fails to account explicitly for the role of expectations in determining the unemployment rate. Here $\hat{E}$ represents expectations with respect to the government's subjective model, and the subscript $t-1$ means that the government updates $\hat{\alpha}_{t|t-1}$ and at each t computes x_{t-1} by solving the time t Phelps problem before observing π_t and u_t. Thus, the government sets policy based on its estimated Phillips curve (4), not the true Phillips curve (1). A self-confirming equilibrium (SCE) is a vector of government beliefs α_{SCE} that is consistent with what it observes in the sense of satisfying the population

[12] If $\text{abs}(\theta_0) > \text{abs}(\theta_1)$, (1) becomes a version of a natural-rate Phillips curve that allows a serially correlated disturbance (Sargent 1999).

least squares orthogonality condition:

$$E\left[\Phi_t(u_t - \Phi_t'\alpha_{\text{SCE}})\right] = 0, \tag{5}$$

where the mathematical expectation is evaluated with respect to the probability distribution of u_t, π_t, and x_{t-1} induced by (1), (2), and the decision rule implied by the Phelps problem defined by $\bar{\alpha}$.

Self-confirming equilibrium outcomes agree with the time-consistent Nash equilibrium outcomes in which policymakers set inflation higher than the socially optimal Ramsey level (see Sargent 1999).[13] Nash inflation is

$$\pi^{\text{Nash}} = \pi^* - \lambda(u^* - u^{**})\left[(1+\delta\tau_1)\theta_0 + \delta\theta_1\right]. \tag{6}$$

The larger are $u^* - u^{**}$, θ_0, and θ_1 in absolute value, the higher is the Nash inflation rate compared to the Ramsey rate π^*.

A self-confirming equilibrium is a population concept that restricts beliefs to be time-invariant and that forms a benchmark – and as it can turn out, a limit point – for our model. Unlike an SCE, in our model after every history, the government updates its beliefs. In particular, the government bases $\hat{\alpha}_{t|t-1}$, its mean estimate of the drifting parameter vector α_t, on the observations up to and including time $t-1$ from the following (misspecified) econometric model:

$$u_t = \alpha_t'\Phi_t + \sigma w_t, \tag{7}$$

$$\alpha_t = \alpha_{t-1} + \Lambda_t, \tag{8}$$

where Λ_t, uncorrelated with w_t, is an i.i.d. Gaussian random vector with mean zero and covariance matrix V. Thus, the government believes that the true economy drifts over time. That is why it continually adapts its parameter estimates with non-vanishing weight on new observations. The innovation covariance matrix V governs the perceived volatility of

[13] As explained by Sargent (1999, chapter 3), the gap between the Ramsey and Nash or SCE outcomes for inflation reflects the benefit to the government of being able to commit to a policy.

increments to the parameters, and is a key component of the model. The mean estimate of α_t for the econometric model (7)-(8) is

$$\hat{\alpha}_{t|t-1} \equiv E(\alpha_t|\mathscr{I}_{t-1}),$$
$$\mathscr{I}_t \equiv \{u_1, \pi_1, \ldots, u_t, \pi_t\}.$$

Let

$$P_{t|t-1} \equiv Var(\alpha_t|\mathscr{I}_{t-1}).$$

Given the government's model, the mean estimates are optimally updated via the special case of Bayes rule known as the Kalman filter. Given $\hat{\alpha}_{1|0}$ and $P_{1|0}$, the Kalman filter algorithm updates $\hat{\alpha}_{t|t-1}$ with the following formula:[14]

$$\hat{\alpha}_{t+1|t} = \hat{\alpha}_{t|t-1} + \frac{P_{t|t-1}\Phi_t(u_t - \Phi_t'\hat{\alpha}_{t|t-1})}{\sigma^2 + \Phi_t' P_{t|t-1}\Phi_t}, \tag{9}$$

$$P_{t+1|t} = P_{t|t-1} - \frac{P_{t|t-1}\Phi_t\Phi_t' P_{t|t-1}}{\sigma^2 + \Phi_t' P_{t|t-1}\Phi_t} + V. \tag{10}$$

An important issue is whether the learning process will converge to a self-confirming equilibrium in which the discrepancy between the government's model and the true DGM vanishes for outcomes that occur thereafter with positive probability. To summarize what we known about this, we scale the innovation covariance matrix as $V = \varepsilon^2\hat{V}$, for $\varepsilon > 0$. Key analytical results from Sargent and Williams (2005) that highlight possible outcomes of the government's learning process are:

(1) In this model, inflation converges much faster to the SCE under Kalman filtering learning than under RLS. The Kalman filter learning rule with drifting coefficients discounts past data more rapidly than the constant gain RLS learning rule.

[14]Many learning models such as Sargent (1999) have focused on a recursive least squares learning (RLS) rule that is closely related to the Kalman filter. Sargent and Williams (2005) show that RLS can be approximated by a Kalman filter with V proportional to $\sigma^2 E(\Phi\Phi')^{-1}$.

(2) As the government's prior belief parameter $\varepsilon \to 0$ (at the zero limit there is no time variation in the parameters), inflation converges to the self-confirming equilibrium (SCE) and the mean escape time becomes arbitrarily long.

(3) As the government's prior belief parameter $\sigma \to 0$ (in the zero limit, either there is no variation in the government's regression error or there is arbitrarily large time variation in the drifting parameters), large escapes from an SCE can happen arbitrarily often and nonconvergence is possible.

(4) The covariance matrix V in the government's prior belief about the volatility of the drifting parameters affects the speed of escape. The covariance matrix V combined with the prior belief parameter ε, affects the speed of convergence to the SCE from a low inflation level.

IV. Empirical Methodology

The theoretical results indicate how very different outcomes can emerge from different government beliefs. The task of this paper is to fit the model to the data and thereby to estimate and quantify the uncertainty about the parameters, σ^2 and V, jointly with the model's other structural parameters, including those governing the "true" expectational Phillips curve (1). Before estimation, we fix the values of δ, λ, π^*, u^{**}, and $\hat{\alpha}_{1|0}$. Group all other free structural parameters as

$$\phi = \{v^*, \theta_0, \theta_1, \tau_1, \zeta_1, \zeta_2, u(C_P), u(C_V)\},$$

where $v^* = u^*(1-\tau_1)$, C_P and C_V are upper triangular such that $P_{1|0} = C_P' C_P$ and $V = C_V' C_V$, and $\zeta_1 = 1/\sigma_1^2$ and $\zeta_2 = 1/\sigma_2^2$ represent the precisions of the corresponding innovations. The notation $u(C_P)$ or $u(C_V)$ means that only the upper triangular part of C_P or C_V are among the free parameters.

The structural parameter $\zeta = 1/\sigma^2$ is not free. It is clear from (9), (10), and (14) that if we scale V and $P_{1|0}$ by κ and ζ by $1/\kappa$, the likelihood value remains the same. There would exist a continuum of maximum likelihood estimates (MLEs) if ζ were not restricted (i.e., the model is unidentified). Some normalization is necessary. Following Sargent and Williams, we impose the restriction $\zeta = \zeta_1$. This normalization implies that the variation that policymakers observe in the unemployment rate is correctly decomposed into variation in the regressors and variation due to exogenous shocks. This assumption has an advantage because it makes limiting results easier to derive.[15]

As we've noted, Sargent and Williams (2005) show that whether monetary policy stays close to a path associated with a self-confirming equilibrium, and when it does not, how it evolves over time, are both very sensitive to the model's parameters (especially the government's belief about the covariance matrix for the drifting coefficients). This sensitivity is what enables us sharply to estimate key structural parameters, including the elements of V.

To take into account parameter uncertainty, we employ the Bayesian method and develop a Monte Carlo Markov Chain (MCMC) algorithm that breaks ϕ into three separate blocks: θ, $\{\zeta_1, \zeta_2\}$, and φ where

$$\theta = \begin{bmatrix} v^* \\ \theta_0 \\ \theta_1 \\ \tau_1 \end{bmatrix},$$

and $\varphi = \{u(C_P), u(C_V)\}$. The prior pdf of ϕ can be factored as:

$$p(\phi) = p(\theta)\, p(\varphi)\, p(\zeta_1, \zeta_2).$$

The prior distributions of both θ and φ take the Gaussian form:

$$p(\theta) = \text{Normal}\,(\bar{\theta}, \bar{\Sigma}_\theta); \tag{11}$$

$$p(\varphi) = \text{Normal}\,(\bar{\varphi}, \bar{\Sigma}_\varphi). \tag{12}$$

[15]Note that an SCE requires the orthogonality conditions, not necessarily the equality restriction $\zeta = \zeta_1$. Indeed, the examples of Sims (1988) allow $\zeta \neq \zeta_1$.

The prior probability density for the precision parameters ζ_1 and ζ_2 is a Gamma distribution:

$$p(\zeta_1, \zeta_2) = \text{Gamma}\left(\bar{\alpha}, \bar{\beta}\right) = \prod_{i=1}^{2} \frac{1}{\Gamma(\bar{\alpha})\bar{\beta}^{\bar{\alpha}}} \zeta_i^{\bar{\alpha}-1} e^{-\frac{\zeta_i}{\bar{\beta}}}. \tag{13}$$

From equations (1) and (2) one can see that the Jacobian transformation from w_{1t} and w_{2t} to u_t and π_t is equal to 1. It follows that the likelihood function is:

$$\mathscr{L}\left(\mathscr{I}_T|\phi\right) = \frac{\zeta_1^{T/2}\zeta_2^{T/2}}{(2\pi)^{T/2}} \exp\left\{-\frac{1}{2}\sum_{t=1}^{T}\left[\zeta_1 z_{1t}^2 + \zeta_2 z_{2t}^2\right]\right\}, \tag{14}$$

where z_{1t} and z_{2t} are the functions of θ and φ:

$$z_{1t} = u_t - u^* - \theta_0(\pi_t - x_{t-1}) - \theta_1(\pi_{t-1} - x_{t-2}) - \tau_1(u_{t-1} - u^*),$$

$$z_{2t} = \pi_t - x_{t-1},$$

where the optimal decision rule depends on φ.

The posterior pdf of ϕ is proportional to the product of the likelihood (14) and the prior $p(\phi)$:

$$p(\phi|\mathscr{I}_T) \propto \mathscr{L}\left(\mathscr{I}_T|\phi\right)\, p(\phi). \tag{15}$$

The posterior distribution of ϕ can be simulated by alternately sampling from the conditional posterior distributions (a Gibbs sampler):

$$p(\theta \mid \mathscr{I}_T, \zeta_1, \zeta_2, \varphi),$$

$$p(\zeta_1, \zeta_2 \mid \mathscr{I}_T, \theta, \varphi),$$

$$p(\varphi \mid \mathscr{I}_T, \theta, \zeta_1, \zeta_2).$$

Appendix C tells how to sample from each of these conditional distributions.

V. Reverse Engineering Estimation

In this section, we present our results. Using the monthly US data described in Appendix A and the prior specified in Appendix B, we estimate ϕ by maximizing the posterior density function. We obtained similar results using maximum likelihood, but the prior is crucial for small sample inference. In estimation, we set $\delta = 0.9936, \lambda = 1, \pi^* = 2$, and $u^{**} = 1$. The value of u^{**} is set at a value low enough to allow Nash inflation to be higher than Ramsey inflation.[16] Setting the unemployment target closer to the natural rate has no effect on our main results.[17]

We set the initial belief $\hat{\alpha}_{1|0}$ at the regression estimate obtained from the presample data from January 1948 to December 1959.[18] We tried to fix $P_{1|0}$ at the value that scales up and down the presample regression estimate $\hat{\sigma}^2(\Phi'\Phi)^{-1}$, but the fit was bad. Similarly, fixing V at the value estimated from a presample-estimated covariance matrix with different scales does not improve the poor fit.[19]Departing from Sargent (1999), therefore, we estimate the government's prior beliefs $P_{1|0}$ and V within the sample. Our MCMC or maximum likelihood algorithm reverse engineers the empirical Phillips curve at each date that, in conjunction with the Phelps problem, rationalizes that date's inflation rate. Estimating $P_{1|0}$ and V gives us the flexibility to succeed in this reverse engineering. Moreover, this flexibility is arguably reasonable. We take the view that the presample data are informative about the government's subjective point estimates (which we fix), but that they substantially understate the government's subjective uncertainty about coefficient innovation volatility V

---

[16]Alan Blinder (1998) emphasizes that the source of time inconsistency in Kydland and Prescott's (1977) Phillips curve example is their specification that $u^{**} \neq u^*$ in the monetary authority's preferences (3). His experience as Vice Chairman of the Federal Reserve let Blinder to question whether the FOMC perceived there to be much of a gap between u^{**} and u^*.

[17]Narrowing the gap between u^* and u^{**} slightly increases the magnitude of the estimated Phillips curve slope parameters θ_0 and θ_1. But this change is small enough as to be essentially inconsequential for our results.

[18]In an earlier draft, we followed Chung (1990) and estimated this belief from the sample data. Since it is influenced by the updated beliefs in the sample, the value estimated this way is as difficult to interpret as that in Chung.

[19]This point is illustrated further in Section VI.2.

(which we estimate). Thus, we use the presample data to pin down the mean of the government's estimate of the empirical Phillips curve, but *not* to estimate the belief innovation covariance matrix V.

We report the posterior estimate of ϕ (evaluated at the peak of the posterior pdf) in Table 1, along with the 68% and 90% probability intervals around the estimate.[20] In our estimation and inference, the regressor vector in the government's Phillips regression (4) is:

$$\Phi_t = \begin{bmatrix} \pi_t & \pi_{t-1} & u_{t-1} & \pi_{t-2} & u_{t-2} & 1 \end{bmatrix}'.$$

As mentioned, among the parameters that we estimate are those of the expectational Phillips curve (1) that we assume truly governs the data. As can be seen in Table 1, the natural rate of unemployment u^* in equation (1) is estimated to be 6.1 and its probability intervals are wide, consistent with the confidence interval in the statistical model of Staiger, Stock, and Watson (1997). Responses of unemployment to inflation surprises (θ_0 and θ_1) are extremely weak and are statistically insignificant by the probability intervals. This is an important finding for us, partly because it implies from (6) that Nash inflation is close to π^* despite the large difference between u^* and u^{**}. Therefore, outcomes close to those associated with the limit point of the mean dynamics are close to the Ramsey outcome. Unemployment is by itself a persistent series and the persistence is tightly estimated.

It can be seen from the estimates and probability intervals of ζ_1 and ζ_2 that their posterior distribution is tight but skewed downward, especially for ζ_1, whose estimate (evaluated at the peak of the posterior pdf) is outside the 90% interval.

The estimated $P_{1|0}$ shows strong correlations (at least above 0.95) among all the elements. The relatively large variance for the drifting coefficient on π_{t-2} (the 4th element) implies that the government is quite uncertain about this coefficient, which affects the uncertainty about other coefficients even though their marginal variances are relatively small.

[20]All probability intervals are derived from the empirical joint posterior distribution generated from a sequence of 50,000 MCMC draws.

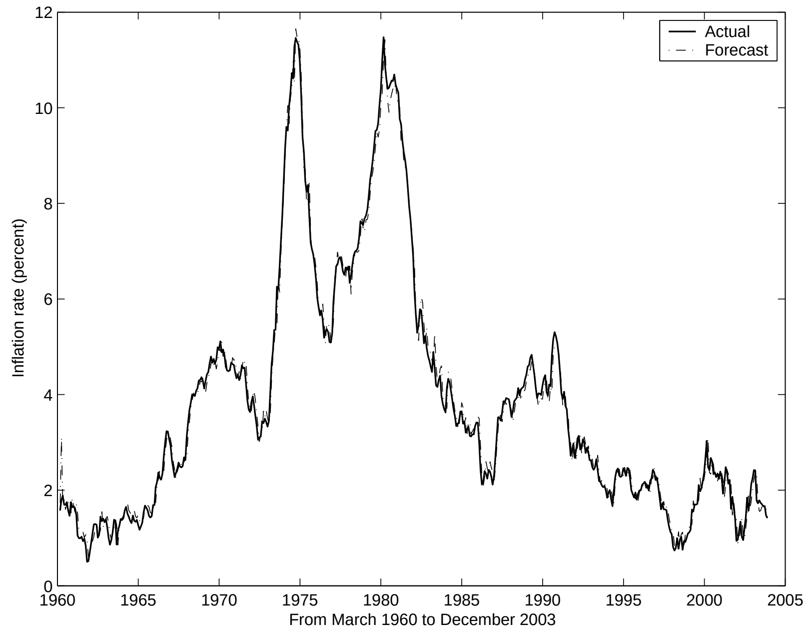

FIGURE 1. Inflation: actual vs one-step forecast (i.e, government controlled inflation)

The estimated V shows strong correlations among the innovations to the coefficients on current and lagged inflation variables. The variances on the inflation variables are large in the sense of Sargent and Williams (2005), implying that the government is willing to adjust its beliefs quickly in response to recent data. The constant term has a very large variance too and affects the coefficients on the lagged unemployment variables because of high correlations, but it has a small influence on the inflation coefficients in the government's model. Because V is not small, the government's beliefs are likely to drift significantly and inflation is likely to escape to the near-Ramsey region. Our estimates of the true expectational Phillips curve (1) imply a negligible difference between the SCE and π^*. We show in Section VII.2 that even when we artificially alter the parameters of (1) to allow the SCE inflation rate to be considerably higher than the Ramsey rate, this large V permits frequent escapes to low inflation rates.

The inflation path produced by the government's inflation policy is plotted against the actual path in Figure 1, and one-step forecasts of unemployment are plotted against the

TABLE 1. Posterior estimates of model parameters

Maximum log value of likelihood (multiplied by prior): 564.92

Estimates of coefficients in true Phillips curve and inflation process
with 68% and 90% probability intervals in parentheses

u^* :	6.1104	(5.2500, 7.1579)	(4.2238, 9.0586)
θ_0 :	−0.0008	(−0.0237, 0.0475)	(−0.0458, 0.0719)
θ_1 :	−0.0122	(−0.0375, 0.0297)	(−0.0589, 0.0526)
τ_1 :	0.9892	(0.9852, 0.9960)	(0.9817, 0.9996)
ζ_1 :	35.6538	(28.7565, 32.4947)	(27.6017, 33.7890)
ζ_2 :	18.97671	(15.6565, 18.2557)	(14.7008, 19.1196)

Estimate of $P_{1|0}$:

10.8705	14.3236	2.2518	−25.4037	−0.9279	−10.1548
14.3236	19.3721	2.9624	−33.9832	−1.1883	−13.5923
2.2518	2.9624	0.4690	−5.2629	−0.1928	−2.1050
−25.4037	−33.9832	−5.2629	59.8997	2.1339	23.9551
−0.9279	−1.1883	−0.1928	2.1339	0.0816	0.8526
−10.1548	−13.5923	−2.1050	23.9551	0.8526	9.5810

Estimate of V:

8.2323	−7.7781	0.9208	4.9782	−0.8136	−41.414
−7.7781	8.1400	0.0303	−5.089	1.9353	68.591
0.9208	0.0303	2.9854	0.1187	3.7012	72.067
4.9782	−5.089	0.1187	3.2032	−1.0548	−39.963
−0.8136	1.9353	3.7012	−1.0548	5.1362	100.6400
−41.414	68.591	72.067	−39.963	100.6400	2588.3000

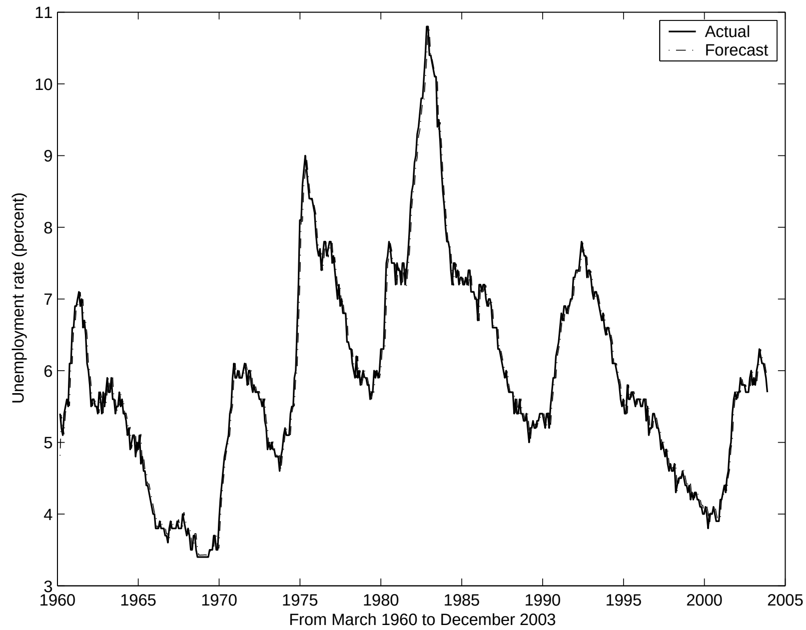

FIGURE 2. Unemployment rate: actual vs one-step forecast

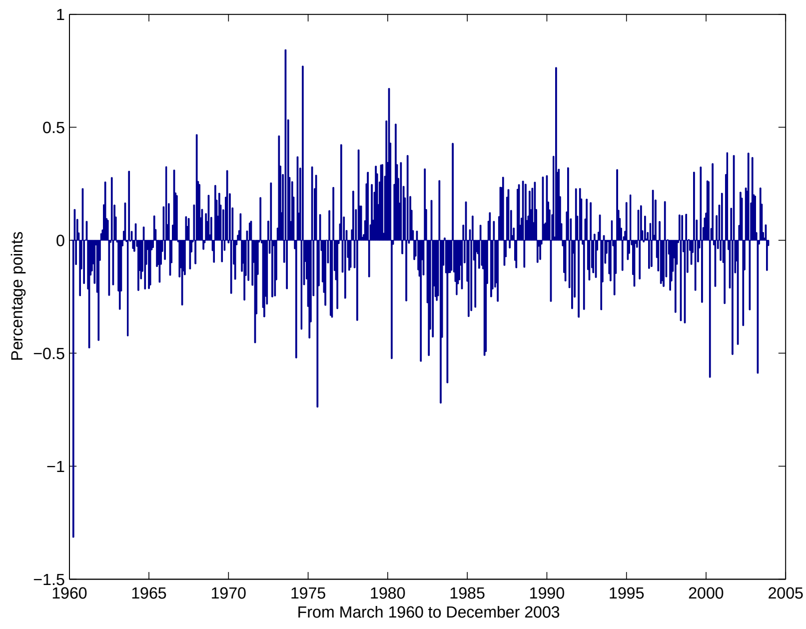

FIGURE 3. Differences between actual values and one-step forecasts of inflation

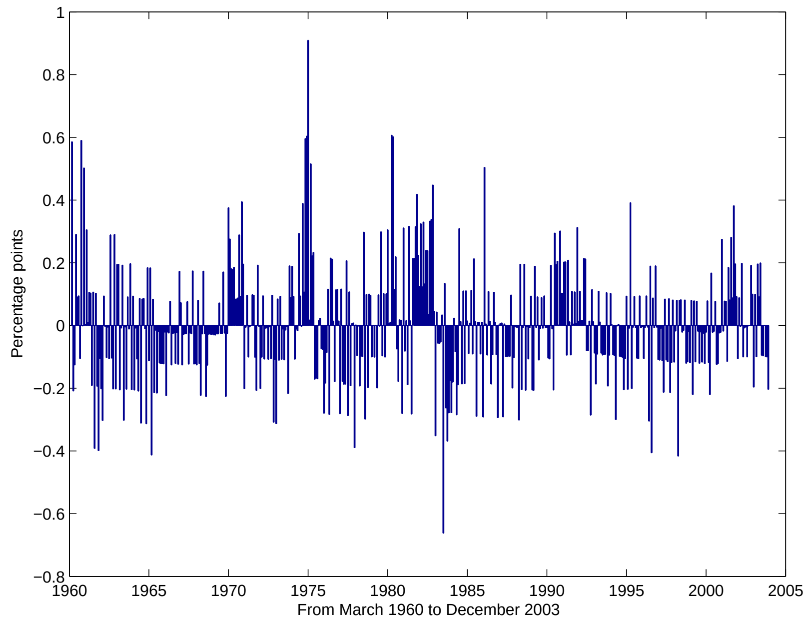

FIGURE 4. Differences between actual values and one-step forecasts of unemployment

actual path in Figure 2.[21] It is evident from these figures that the model explains the low-frequency movements of inflation well; so well, in fact, that it is difficult to discern the difference between the series.[22] By this fit criterion, our reverse engineering exercise is a success, especially compared to those carried out by Chung (1990) and Sargent (1999). Figures 3 and 4 plot the one-step forecast errors for inflation and unemployment, showing that for most of the sample, the forecasts are within one half a percentage point of the realized value.

These forecast errors are comparable to those from BVAR models with the standard prior settings proposed by Sims and Zha (1998). The root mean square error (RMSE) and the mean absolute error (MAE) are 0.225 and 0.179 for our model, 0.397 and 0.297

[21]Inflation policy (x_{t-1}) is sharply estimated. Although we do not plot the error bands around the estimated x_{t-1} to avoid visually clustering Figure 1, they are quite tight and track the rise and fall of actual inflation well.

[22]These empirical results provide a formal justification for Ireland (2005)'s key assumption that persistent changes occurred in the Federal Reserve's inflation target in the 1970s.

for the BVAR with one lag (BVAR(1)), and 0.357 and 0.272 for the BVAR with 13 lags (BVAR(13)).[23]

We also use the Schwarz criterion (SC) to compare maximum log values of the likelihood multiplied by the prior among our learning model, the BVAR(1), and the BVAR(13).[24] The SC value is 564.92 for our model, 313.98 for the BVAR(1), and 309.37 for the BVAR(13). Our learning model appears to dominate the two atheoretical models.

To see whether these asymptotic results holds in finite samples, we compute the marginal data density (MDD) for our learning model, using the modified harmonic mean method described in Geweke (1999) and Propositions 1 and 2 in Appendix C. The log MDD value is 424.75.[25] In comparison, the log MDD value is 172.05 for the BVAR(1) and 244.65 for the BVAR(13).[26] As measured by Bayes factors (which would put essentially zero weight on the BVARs), our learning model again dominates the BVARs.[27]

---

[23]We also consider the Keynesian Phillips curve in the government problem by putting inflation on the left side of the regression, as discussed in King and Watson (1994)). The resulting RMSE is 0.478 and MAE is 0.332, substantially bigger than our model's. The poor fit of the Keynesian model is consistent with the findings of Cogley and Sargent (2005a).

[24]The Schwarz criterion (sometime called Bayesian information criterion) is used to adjust log likelihood by the number of degrees of freedom times log of sample size divided by 2. This criterion is very useful because it can be readily computed from the estimates reported in Table 1 and because, under standard regularity conditions, it approximates the posterior odds well if the sample is large enough. We follow Sims and Zha (2004) and use the likelihood multiplied by the prior instead of likelihood itself, because models with a large number of parameters are better characterized by the likelihood multiplied by a prior. The same asymptotic reasoning that justifies the Schwarz criterion based on the likelihood applies to the likelihood multiplied by a prior.

[25]The computed value is based on one million posterior draws, which take six days on a Pentium-IV PC desktop. Earlier drafts of this paper report a much lower value that was incorrect because we neglected the constant terms in the likelihood function and the pdf weight when using the modified harmonic mean method.

[26]These results differ from the SC results. It is well known, however, that the Schwarz criterion tends to favor VAR models with shorter lags.

[27]The MDD values may be sensitive to the priors. But the differences between the MDD values of our learning model and the BVAR models are large enough for us to conclude that our model appears to dominate.

Higher Bayes factors, however, do not necessarily imply that our learning model outperforms BVARs in predicting the rise and fall of inflation. In Sections VI.3 and VI.5, therefore, we compare the performances of both our learning model and BVARs in forecasting *longer-term* inflation . There we show that for forecasting low-frequency movements in inflation, our model performs as well as or better than the BVARs. Without any assumption about exogenous components of the persistence of inflation, the government's inflation policy explains, almost entirely, the rise and fall of post-war American inflation (Figure 1). This result had not been achieved in previous work (e.g., Sims 1988, Chung 1990, and Sargent 1999).

VI. Further Empirical Results

VI.1. **Shocks and Beliefs.** In our model, the rise and fall of inflation is driven by the Phelps problem in conjunction with the government's belief in an exploitable tradeoff between inflation and unemployment, which leads to a high inflation rate in the early 70s. But then occasional sequences of stochastic shocks lead the government temporarily to believe that it can cut inflation with no rise in unemployment, which leads to rapid disinflations in the early 80s. During these episodes, the government learns a version of the natural rate theory in which the sum of the coefficients on inflation is nearly zero in its model, reflecting a vertical long-run Phillips curve.

The evolution of the government's updated beliefs is displayed in Figure 5. The sum of the coefficients on inflation becomes very negative in the early 70s and stays quite negative until the late 70s. Although the sum of the inflation coefficients is still negative, in the 1980s it is small enough to induce policymakers to decide to cut inflation without worrying much about costs in unemployment.

Figure 6 displays the subjective covariations in the drift innovations of some key functions of parameters in the government's Phillips curve, derived from our estimated V reported in Table 1. These key parameters are the sum of the coefficients on current and lagged inflation variables ($\alpha_1 + \alpha_2 + \alpha_4$), the sum of the coefficients on current and lagged

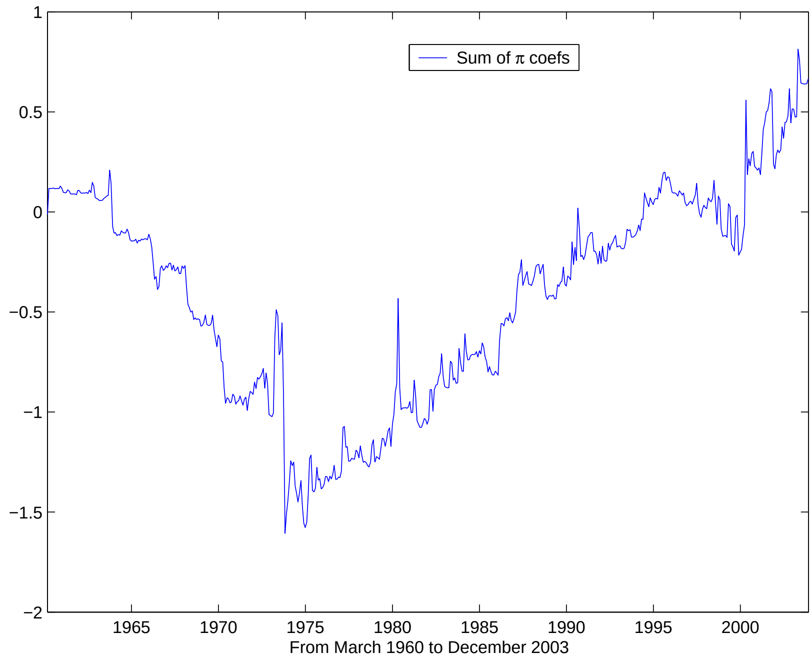

FIGURE 5. Evolution of the government's beliefs

unemployment variables ($1-\alpha_3-\alpha_5$), and the coefficient on the constant term (α_6). As shown by the symbol "$*$" in the first row of graphs of Figure 6, the estimated constant coefficient has a large, positive value while the sum of the estimated inflation coefficients is quite negative. This combination leads to a high perceived tradeoff between unemployment and inflation in December 1973.

In contrast, at the point associated with the SCE (indicated by the symbol "$\circ$" in the second row of Figure 6), the estimated constant coefficient is small and the sum of the inflation coefficients is near zero, providing the government no incentive to inflate in pursuit of lower unemployment.

The probability ellipses shown in Figure 6 are quite large along the dimension of the constant coefficient. The large variation implies that a tradeoff between inflation and unemployment can be severe if there is a high probability that the constant coefficient and the sum of the inflation coefficients fall far within the north-west quadrant, as in the case of the

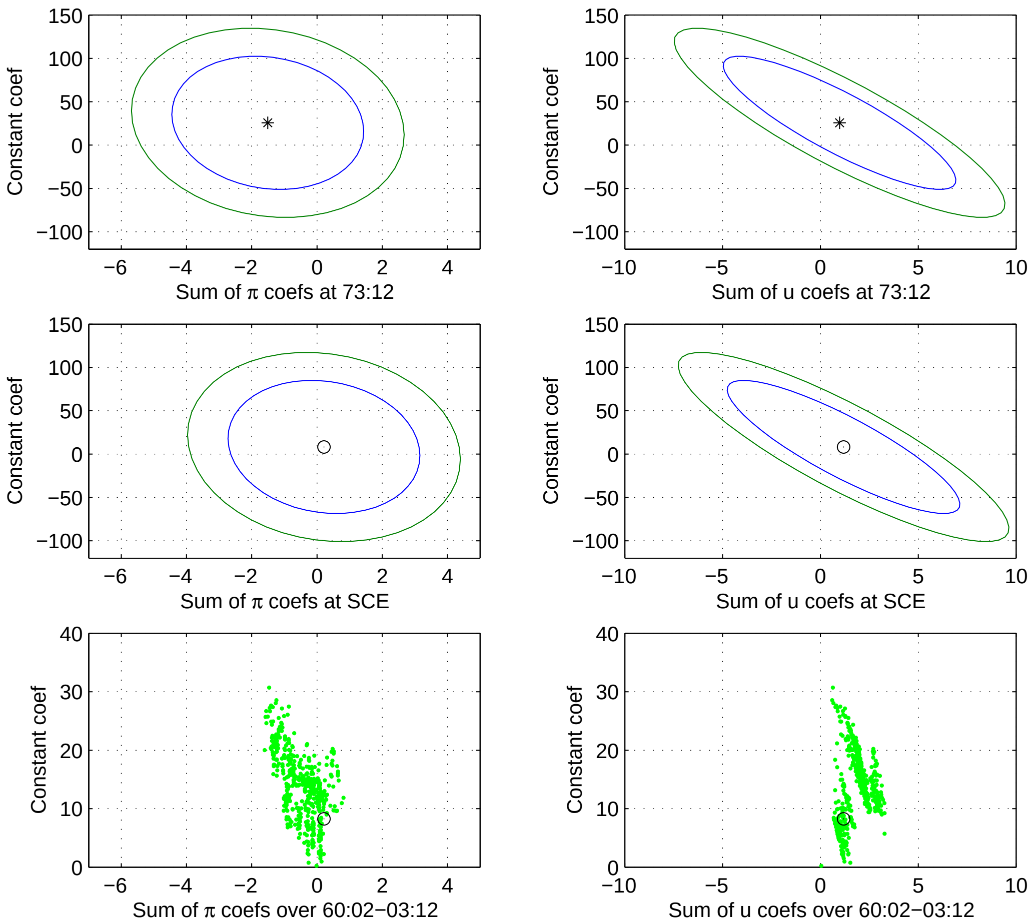

FIGURE 6. 68% and 90% probability ellipses about key parameters in the government's Phillips curve. The first row is based on the observation at 73:12; the second row is based on a limiting case associated with an SCE; the third row displays scatter plots of the estimates throughout our 60:02-03:12 sample. The asterisk symbol * in the first row depicts the government's estimates at 73:12. The circle symbol $\circ$ in the second and third rows depicts SCE values, which also equal limiting estimates from the mean dynamics.

upper-left graph. The bottom-left graph shows the historical estimates of these two belief

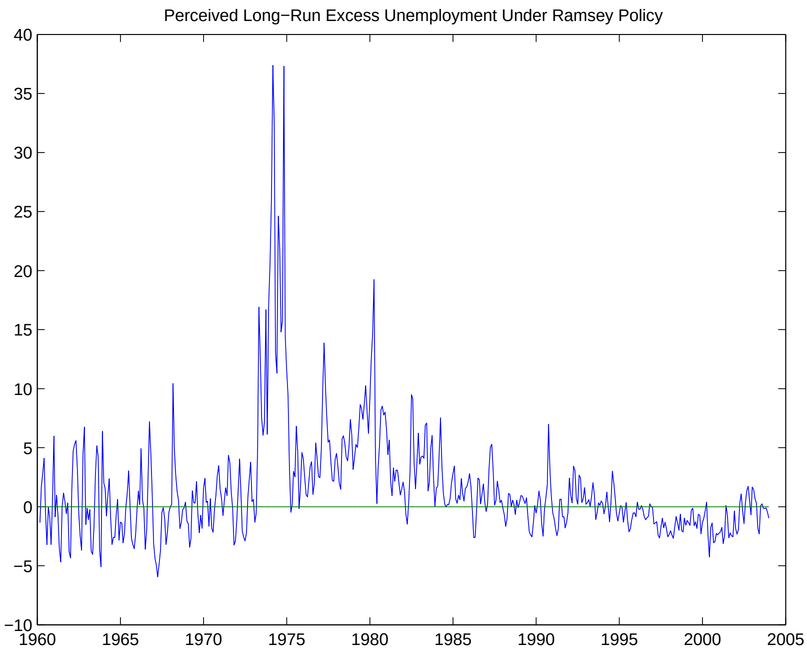

FIGURE 7. Perceived long-run excess unemployment under the Ramsey policy of 2% inflation according to the government's model.

parameters, induced by the particular sequence of shocks throughout our post-war sample. The area in which the sum of the inflation coefficients is less than -1 and the constant coefficient is greater than 15 covers most of the estimates for the 70s.

The constant and the sum of the unemployment coefficients are highly but negatively correlated, as shown in the first two graphs in the second column of Figure 6. Later we will see that in the transition to the SCE, the economy may go through periods of very volatile inflation. If $1-\alpha_3-\alpha_5$ and α_6 frequently have opposite signs because they are negatively correlated, the value of (16) would often be negative because of a large value of α_6. This in turn would prompt the Phelps planner to deflate a lot to stabilize his objective function, thereby causing volatile fluctuations of inflation. These volatile outcomes occur when these two parameters fall in the south-east and north-west quadrants. Fortunately for US inflation outcomes, our historical estimates have been concentrated around the north-east quadrant, as shown in the bottom-right graph. It is only in out of sample simulations that we enter the more volatile regions. Exposure to those out of sample possibilities is a byproduct of the large V that, in conjunction with the Phelps problem, Bayes' rule prompts us to use to reverse engineer the government's choice of inflation.

The three belief parameters discussed above are key inputs to the government's perceived sacrifice ratio. The government's model makes it expect that if it implements the Ramsey policy π^* of 2% inflation,[28] then over the long term, unemployment will on average exceed the natural rate u^* by the amount:

$$\frac{\pi^*(\alpha_1+\alpha_2+\alpha_4)+\alpha_6}{1-\alpha_3-\alpha_5}-u^*. \tag{16}$$

Under any inflation policy in history that differs from Ramsey, the tradeoff will be a simple scaled version of (16) in proportion to that difference.[29] Figure 7 plots the quantity (16). Here we see that, throughout the 1970s, the government's model implied that substantial increases in unemployment would result from a low inflation policy.[30]It wasn't until the early 1980s that this ratio fell nearly to zero, at which time the disinflation commenced. This point will be reiterated in Sections VI.5 and VI.6 where we present longer-term forecasts and counterfactual paths around that time.

VI.2. **Importance of Cross-Equation Restrictions.** As we've already noted, the flexibility that a large V gives our model is crucial for giving us the ability to reverse engineer government beliefs that, intermediated by the Phelps problem, account for the government's decisions about the predictable part of inflation x_{t-1}. In particular, our findings tell us to

---

[28]This is the Ramsey policy, i.e., the optimal policy under commitment, given knowledge of the true DGM.

[29]Note that our measure of the sacrifice ratio differs from the more conventional usage, which gives the cost of disinflating from a current inflation rate. Instead, ours is a long-run measure, independent of current inflation.

[30]A temporary drop in this sacrifice ratio around 1976 led to a temporary decline in inflation around that time. See Cogley and Sargent (2005a) for a story in which the government was deterred from stabilizing in the mid 1970s because it attached a small positive probability to a model that assigned high unemployment costs to a rapid disinflation.

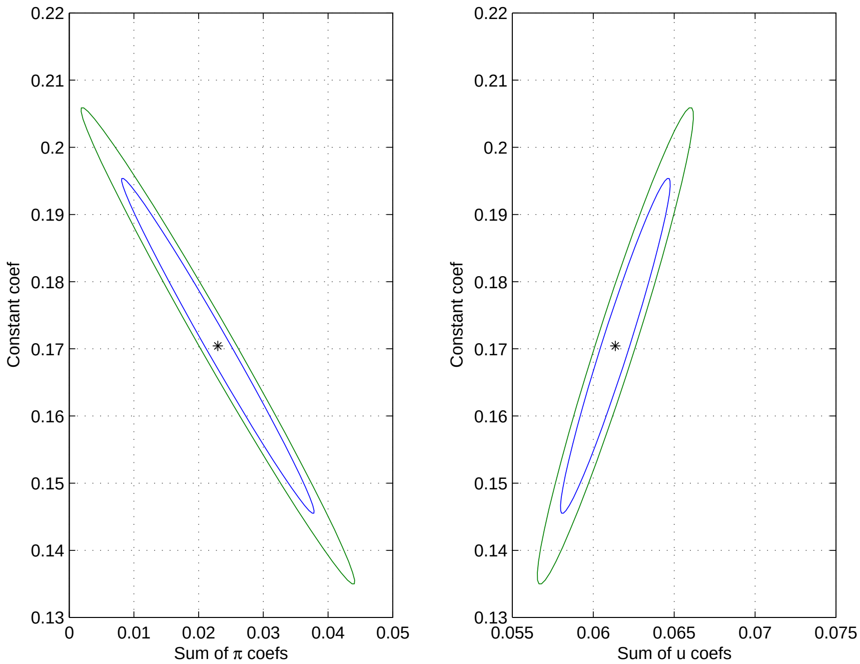

FIGURE 8. 68% and 90% probability ellipses about key parameters in the government's Phillips curve, derived from the estimated V without imposing cross-equation restrictions. The asterisks mark the estimates of these belief parameters at 73:12 (when inflation was quite high).

attribute the empirical failure of previous work with similar models by Chung, and Sargent to the fact that they assumed a particular form for the key matrix V in (8) that governs the innovations to the parameters in the government's model.

To highlight the importance of V, we can estimate V (and $P_{1|0}$) directly with (7) and (8), thereby abstaining from imposing the Phelps problem. These estimates can serve as a benchmark for what impacts on V occur from our imposing the cross-equation restrictions via the Phelps problem discussed in Section VI.1.

Figure 8 displays the covariations in the key belief parameters when the restrictions from the Phelps problem are *not* imposed. Compared to Figure 6, where the restrictions from the Phelps problem are imposed, the ellipses in Figure 8 are very tight – so tight that if we were to combine Figure 8 and the first row of graphs in Figure 6, the tight Figure 8 ellipses would appear as short thin lines. Furthermore, the SCE values are far outside the Figure 8 ellipses.

We have already discussed theoretical reasons that make the V matrix so important and how different specifications of it affect the speed, direction, and stability of the learning dynamics. The V depicted in Figure 8 and those imposed by Chung and Sargent are substantially smaller than what we estimate when we impose the Phelps problem. Relative to the estimate of V that we get when we impose the cross-equation restrictions coming from the Phelps problem, both Chung and the Figure 8 calculations constrain how learning could occur, and diminish the variation in the data that can be explained by evolving government beliefs.

Figure 9 shows what happens when we reestimate the model in the fashion of Chung and Sargent, imposing our estimate of V from Figure 8. The fit deteriorates substantially. The government's optimal policy completely misses the two peaks in inflation in the 1970s, which is what Sargent (1999) found.[31] Chung (1990) and Sargent (1999) found that with their choices of V, the government should have cut inflation much earlier than actually occurred. Our results show how that outcome came from attributing to the government particular beliefs about how its model changes over time.

By imputing to the government the necessary "openness to recent data" that is required by the cross-equation restrictions called for by the Phelps problem, the rise and fall of inflation can be much better explained by the evolution of government's beliefs in response to a particular sequence of shocks in the 70s and 80s. [32]

[31]If we use the sample estimate of the second moment matrix and we choose the proportionality factor so that the new V matrix has the same norm as our estimate, the fit would be as poor as Figure 9. Similarly, if the originally estimated V in Section VI.1 is scaled down by, say, 0.01 so that inflation dynamics are governed by the SCE, the implied inflation policy would completely miss the rise and fall of actual inflation.

[32]One can infer from reading historical records of the Federal Open Market Committee that decision makers spent enormous amounts of time evaluating current economic conditions and that policy deliberations were dominated by interpretations of very recent changes in economic data. Even in the Greenspan era, policymakers' beliefs seemed to be heavily influenced by new developments (see various chapters in

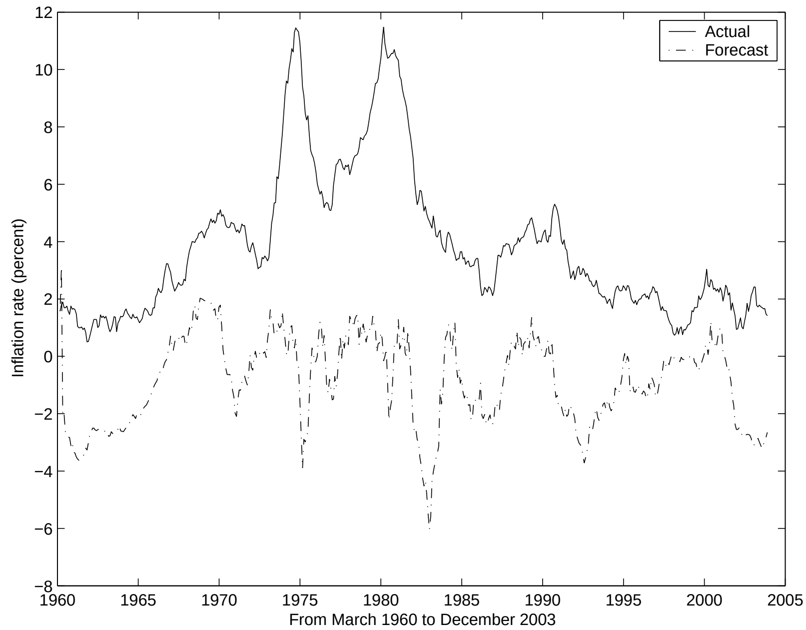

FIGURE 9. Actual inflation and one-step prediction from the benchmark model in which V is estimated without imposing the cross-equation restrictions.

VI.3. **Longer-Horizon Inflation Forecasts.** Longer-term forecasts of inflation play an important part in policy discussions at Federal Open Market Committee (FOMC) meetings. At each FOMC meeting, Federal Reserve economists prepare a report called the Greenbook that forecasts various economic variables over the two-year horizon. How well would our learning model would do in producing two-year-ahead forecasts of inflation throughout the sample, as compared to the BVAR(13) model?

Figure 10 depicts the two-year-ahead median forecasts of inflation from our model and also ones from a BVAR(13).[33] The forecast value and the actual value are so aligned that if

(Chappell, McGregor, and Vermilyea, 2005)). Our reverse-engineering estimate of V quantifies the FOMC's preoccupation with recent data in the context of a *formal* model.

[33]At each time t, we first draw a sequence of structural shocks w_{1t+k} and w_{1t+k} defined in (1) and (2) for $k = 1, \ldots, 24$. Conditioning on the estimated values of the structural parameters, the estimated beliefs at t, and the data $\mathscr{I}_t$ defined in Section III, we then employ (1) and (2) to generate the forecasts u_{t+k} and π_{t+k}

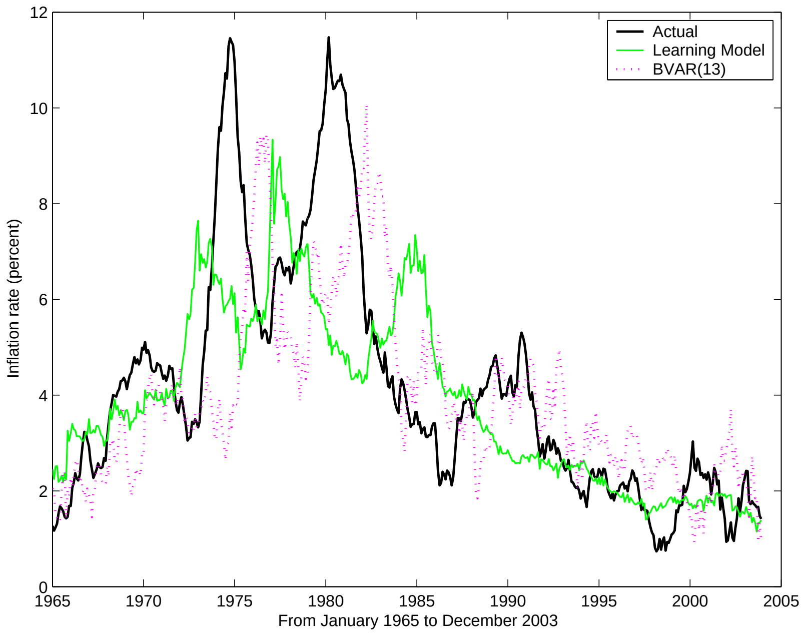

FIGURE 10. Two-year-ahead inflation forecasts: learning model versus BVAR(13)

the forecast were right on the target, these values would be atop one another. As shown in the figure, the learning model produces forecasts that differ substantially from the BVAR.[34] Our learning model predicts the first two rises of inflation about between 1 and 2 years two early and a third small rise of inflation about 5 years too early. And it predicts a permanent fall of inflation after 1985 with less forecasting volatility than the BVAR(13). By contrast, the two-year ahead forecasts of inflation from the BVAR(13) seem to lag the rise and fall of actual inflation.

Figure 11 traces these predicted accelerations of inflation to features of the government's beliefs that lead the Phelps planner to expect to "step on the gas" each of these three times.

---

by recursively solving the inflation policy via the Phelps problem. We repeat this simulation 1000 times and calculate the median of all simulated values of π_{t+24}. This computation takes about 40 hours on a Pentium-IV PC desktop.

[34]The RMSE and MAE are 1.9292 and 1.3233 for the learning model, 2.3939 and 1.6809 for the BVAR(1), and 2.0861 and 1.4617 for the BVAR(13).

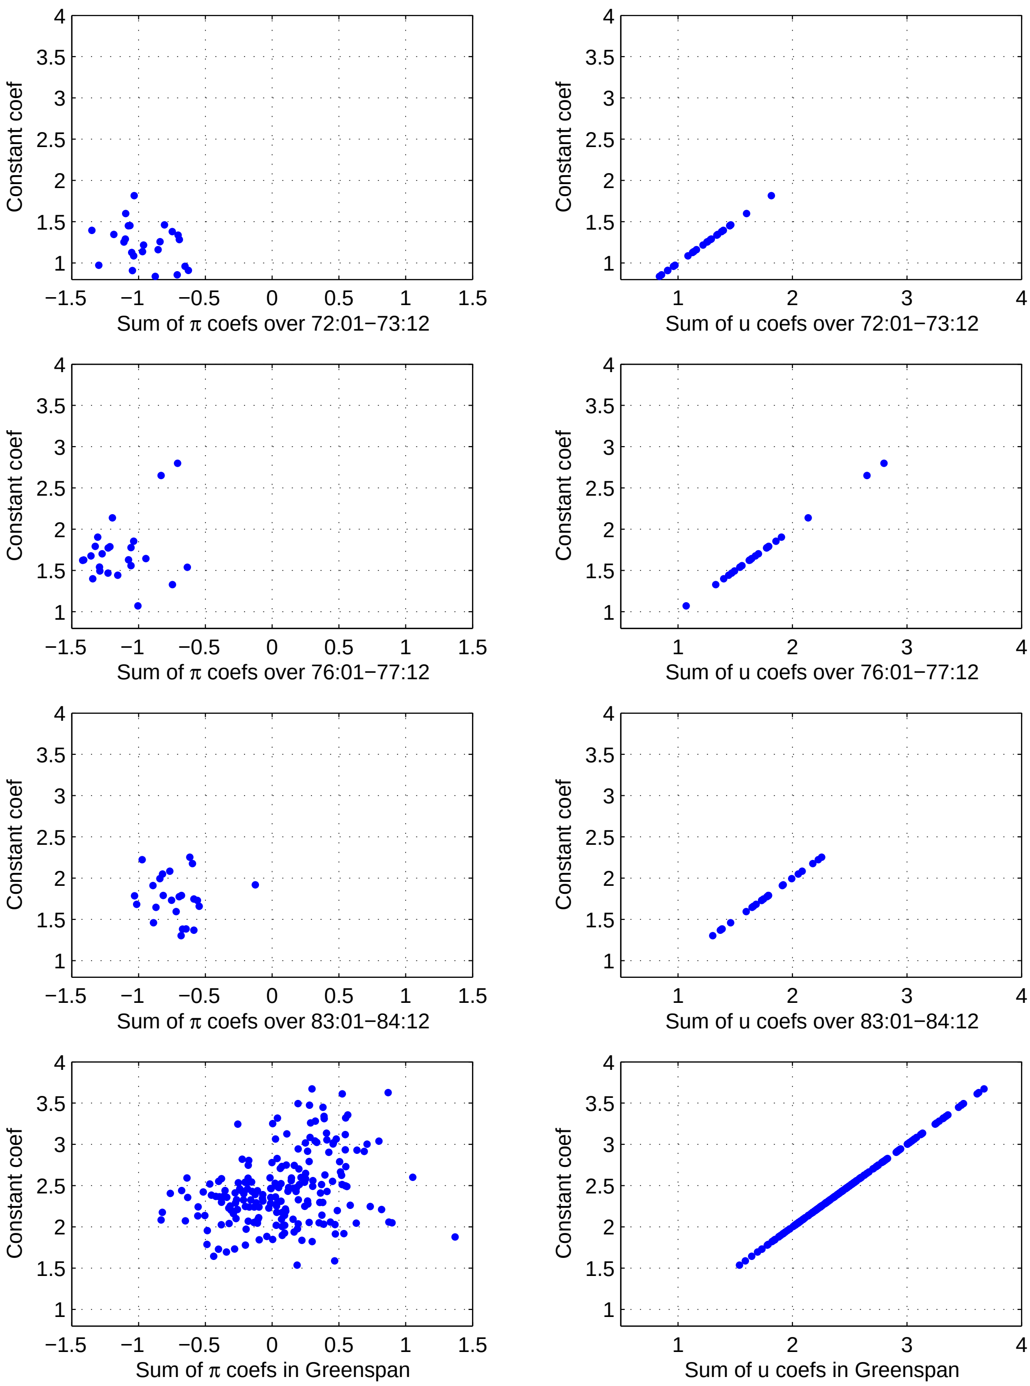

FIGURE 11. Estimates of key step-on-gas parameters over the three predicted run-up periods and over the Greenspan era. The first row shows the evolution of these belief parameters for 72:01-73:12 (the first predicted run-up period); the second row for 76:01-77:12 (the second predicted run-up period); the third row for 83:01-84:12 (the third predicted run-up period); and the fourth row for 87:07-03:12 (the Greenspan era).

The left graph in the first row shows that the sum of inflation coefficients move from left to right over time, getting more negative and prompting the government to step on the gas. In the right graph, one can see that the sum of unemployment coefficients and the constant coefficient are in the north-east quadrant, indicating that the inflation forecast is stable for this period, as discussed in Section VI.1. The second and third rows show similar patterns, with differences in how negative the sum of inflation coefficients gets over time. The left graph in the fourth row, however, reveals a completely different story. The sum of inflation coefficients moves toward zero over time and then passes into positive territory. Thus, the government faces at most weak inflation-unemployment trade-offs. These results explain why, after a third run-up of actual inflation between 1986 to 1990, the government would not want to step on the gas. Interestingly, the BVAR continues to predict a run-up even after 1990.

VI.4. **Good low frequency outcomes.** Figure 12 displays the four-year-ahead predictions from our model and the BVAR(13). Neither model predicts the magnitude of the rises of inflation that occurred. But our model captures the timings of the first two rises almost perfectly, while the predictions of the BVAR(13) again lag behind. The RMSE and MAE are 1.761 and 1.241 for our model, 2.838 and 2.195 for the BVAR(1), and 2.433 and 1.820 for the BVAR(13). Our model's 4-year forecast errors are smaller than its 2-year forecast errors, while the forecast errors from the BVARs are larger for the 4-year horizon than for the 2-year horizon.

VI.5. **Two Peaks and an Enduring Decline.** To reinforce the results in the last section, we now analyze in further detail how the model forecasts the two peaks of inflation in the 1970s and the sharp decline in the early 1980s. We look at both the point forecasts and the associated distributions at various forecast horizons, conditioning on the estimated values of the structural parameters. We use Monte Carlo simulations to assess the distribution of forecasts going forward over four year horizons from different initial conditions. In each case, we take the estimated beliefs at the starting date and repeat 5000 simulations of 50

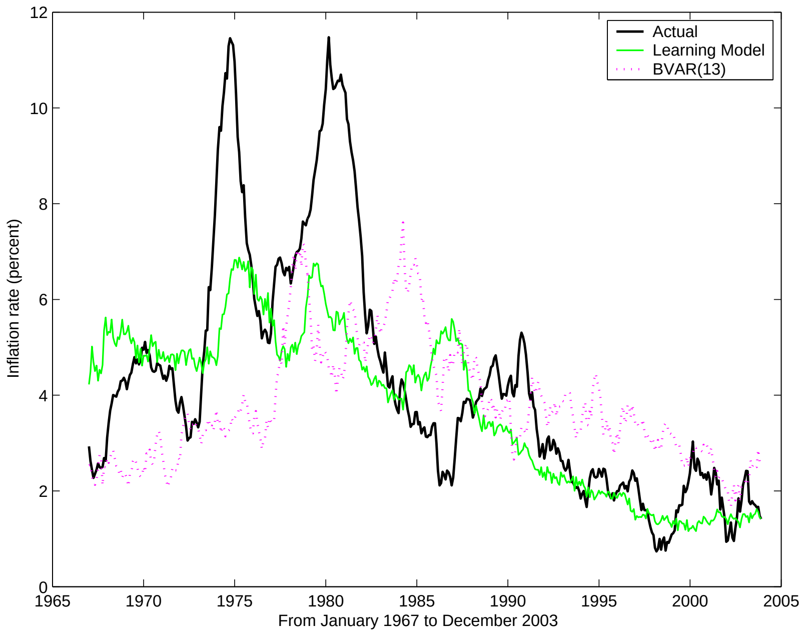

FIGURE 12. Four-year-ahead inflation forecasts: learning model versus BVAR(13)

periods.[35] We then plot the actual experienced inflation and the median forecast along with 68% and 90% probability bands. In each plot, the initial condition is shown as date zero, from which we look forward 50 periods.

Figure 13 reports the forecasts. The upper left panel starts in January 1973 when inflation was at a very low level (3.3%). This is also when we say that the government most overestimated the tradeoff between inflation and unemployment (see Figure 5). According to the model, the government exploited the tradeoff and pushed up inflation to lower unemployment. The model predicts a steadily rising inflation path as high as 10% towards the end of the 4-year horizon (the upper 90% band), and gives little probability to a lower inflation rate in the medium run.

Due to a sequence of shocks, the inflation path reached its peak earlier than the model predicts. But this is a treacherous period in which to predict, and as we show later in this

[35]Adding uncertainty in the parameters would widen our forecast bands only a little.

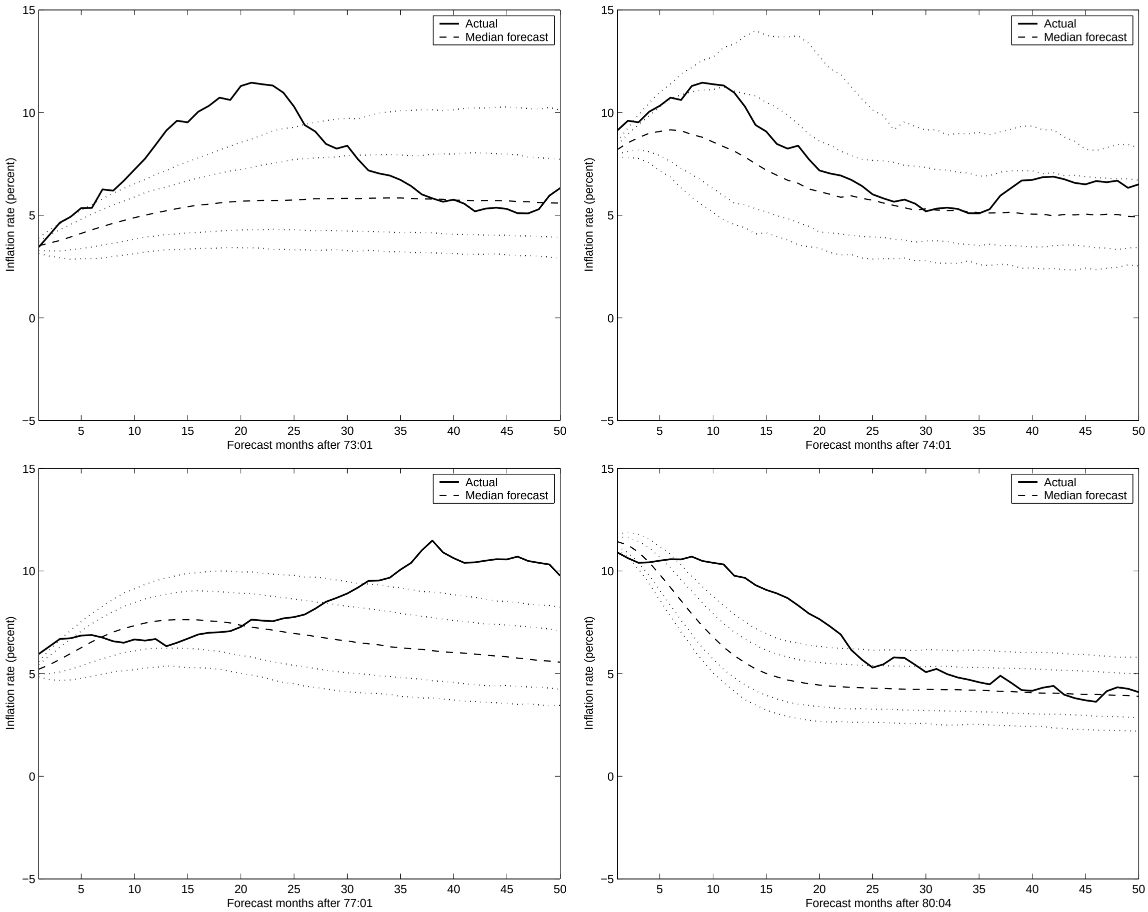

FIGURE 13. Dynamic forecasts of inflation with 68% and 90% error bands, using as initial estimated conditions at 73:01, 74:01, 77:01, and 80:04.

section, our model's prediction of rising inflation compares favorably to predictions coming from alternative statistical models.

A year later in January 1974, which is shown in the upper right panel of Figure 13, inflation had continued upward, now reaching 8.4%. Here we see that the model tracks the actual inflation path quite well, predicting a further increase in inflation prior to a return to lower levels.

January 1977 was another difficult time to predict inflation because inflation was at its trough and a second run-up was about to begin. Although actual inflation reached its peak at a later date, the model assigns an overwhelming probability to higher inflation and the upper 90% reaches as high as 10% (the lower left panel).

The disinflation episode in the early 1980s is often interpreted as reflecting the intellectual triumph of the rational expectations version of the natural rate theory. What does our learning model say about this period? Would the government continue to pursue a higher inflation policy? After all, from the vantage point of April 1980 when inflation reached its second peak, most forecasting models either predict that inflation was very likely to go higher than it actually did, or they fail to predict the fall of inflation. The lower right panel of Figure 13 displays the forecast from our learning model. While actual inflation declines at a somewhat slower speed than the model predicts in 1980 and 1981, the forecast of a fast decline in inflation is remarkable. The model's prediction is especially good further in the forecasting period. Unlike many forecasting models, our model gives almost no probability to rising inflation in the medium horizon, because the tradeoff between inflation and unemployment by then is not high enough for the government to pursue double-digit inflation.

We now compare the model's forecasts with those from the BVARs. Figure 14 displays the forecast made at 73:01 from the BVAR(1). The 68% and 90% error bands are produced by simulating the VAR shocks while holding the parameter estimates fixed at those obtained using the 60:01-03:12 sample, the same procedure as we earlier applied to our learning model. As can be seen, this BVAR fails to predict any rise of inflation with a significant probability. And the upper 90% band is well below 10%.

Figure 15 shows the forecasts of inflation at the various dates from the BVAR(13), resembling Figure 13. The forecasts at 73:01 from the BVAR(13) are not so different from those from the BVAR(1) except the error bands are much wider. They give probability half to a decline of inflation. For the forecasts at 74:01, the BVAR forecasts are comparable to those from our learning model. The forecasts at 77:01 from the BVAR(13) again give probability half to a decline of inflation while the forecasts from our learning model (Figure 13) put a vast majority of probability to rising inflation. For the forecast at 84:04, the

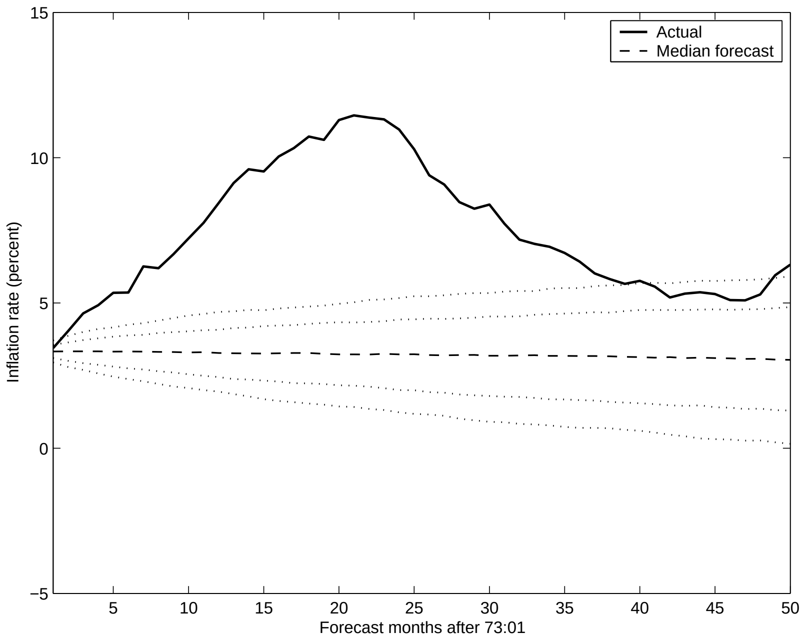

FIGURE 14. Dynamic forecasts of inflation from BVAR(1) at 73:01, with 68% and 90% probability bands.

BVAR(13) predicts a decline of inflation. But our learning model predicts a much sharper decline of inflation with narrow bands while the BVAR assigns considerable probability to higher inflation than the actual path. Overall, our learning model performs as well or better than the BVARs in explaining the rise and fall of inflation at these crucial dates.

VI.6. **Counterfactual Exercises.** As a way to quantify the role of econometric policy evaluation in the government's learning process, we use our estimated classical model to calculate what would have happened if the government's beliefs had differed from our estimates. All of the results in this section condition on estimates of the historical shocks of unemployment and inflation that we infer from our model estimates. We treat these shocks as random and exogenous in our counterfactual exercises.

The first episode begins in 1964:01. As seen from Figure 5, there is still little belief in the inflation-unemployment tradeoff in the early 1964, but by then end of 1973 the sum of

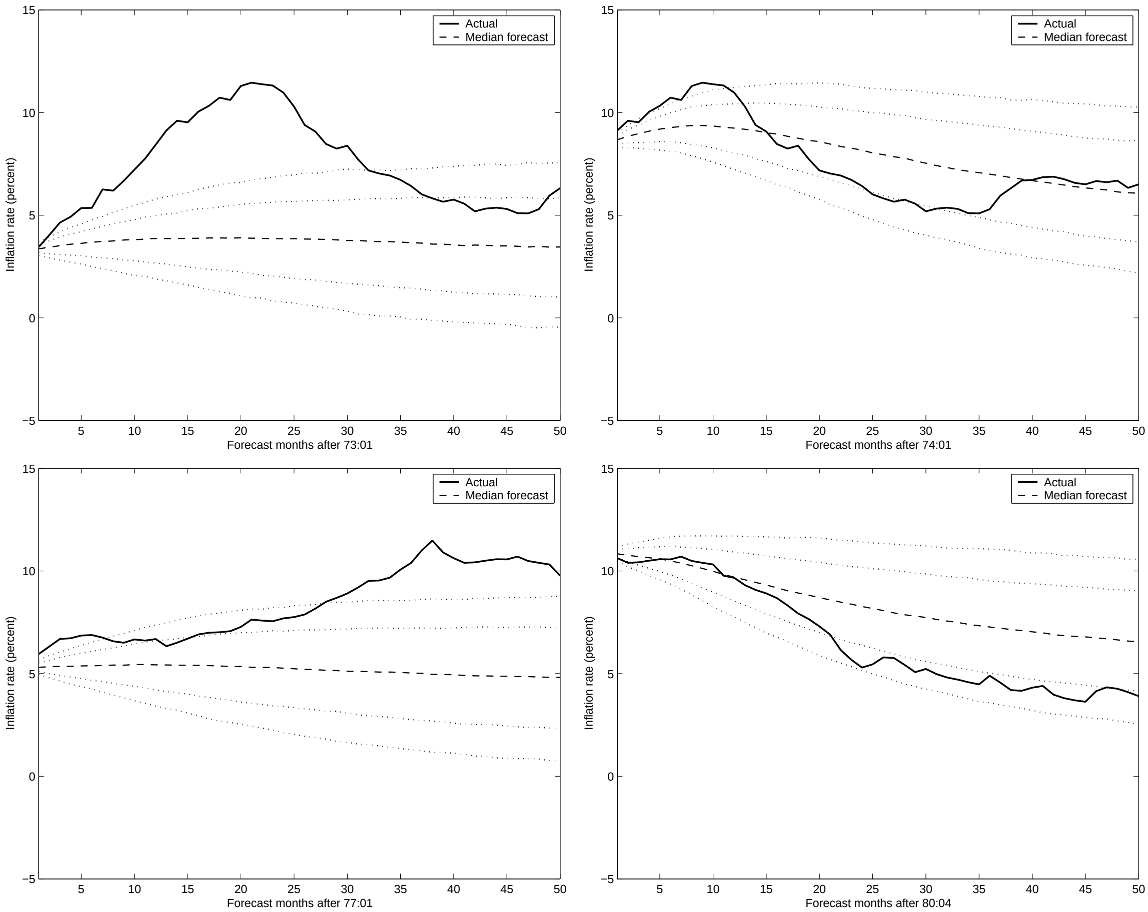

FIGURE 15. Dynamic forecasts of inflation from BVAR(13) at 73:01, 74:01, 77:01, and 80:04, with 68% and 90% error bands.

the inflation coefficients is most negative.[36] Such continual adaptation of beliefs towards a bigger inflation-unemployment tradeoff gives the government an incentive, through the Phelps problem, to run a high inflation policy. This can be seen indirectly in Figure 7 as the perceived costs of low inflation rise dramatically in the early 1970s. To obtain a more direct comparison, suppose the government's beliefs had been frozen at the 64:01 initial condition. As shown in Figure 16, the inflation path would have been smoother and

[36] See Sargent (1999, chapter 5) for how the sum of coefficients on π affects the advice rendered by the Phelps problem.

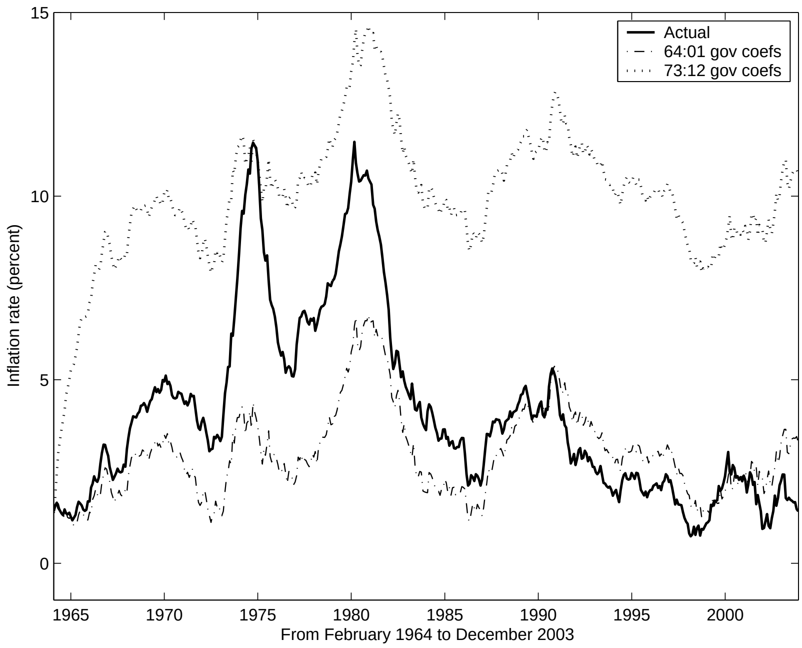

FIGURE 16. Inflation dynamics with fixed beliefs at the dates 64:01 and 73:12.

FIGURE 17. Inflation dynamics with fixed beliefs at the dates 79:10 and 73:12.

avoided much of the two large run-ups of actual inflation in the 70s. To take an opposite example, we replace the government's 64:01 beliefs with the 73:12 beliefs and fix them throughout the history. In this case, Figure 16 shows that inflation would have been much higher than was actually experienced throughout the sample and would have continued to stay around 10%.

Figure 17 displays a second episode beginning with 1979:10, when Volcker's disinflation policy took place. As we've seen, if the government had held fixed to its 1973:12 perceived tradeoff, inflation would have stayed much higher. On the other hand, if the government's belief at 1979:10 had been fixed throughout the rest of the history, inflation would have come down to 5% by 1986 due to the sequence of historical shocks, but there would have been a tendency to return to a higher inflation level. These outcomes show the important role that we assign to adapting government beliefs in the process of achieving lower inflation. With the same sequence of historical shocks, actual inflation came down and remained low as shown by the inflation path in Figure 17. Although the government's beliefs at the end of 1979 favored a disinflation, Figure 5 shows how the government's views continued to evolve to favor a low inflation policy. The experience of disinflation and continued low inflation led the government away from believing in an exploitable Phillips curve tradeoff.

These exercises suggest that while the rise of inflation in the 70s was caused by the government's misperceiving the Phillips-curve relationship, the fall of inflation in the 1980s can be explained by an econometric policy evaluation procedure that embodies adaptive beliefs.[37]

[37]Changes in beliefs do not necessarily imply changes in the linearized policy rule in which x_t is regressed on its own lagged values and current and lagged unemployment variables. In our case, because x_t tracks the actual inflation path so well, our results are consistent with reduced-form empirical findings that changes in the policy rule or the inflation process are difficult to detect statistically (Cogley and Sargent 2005b, Primiceri in press, and Sims and Zha 2004).

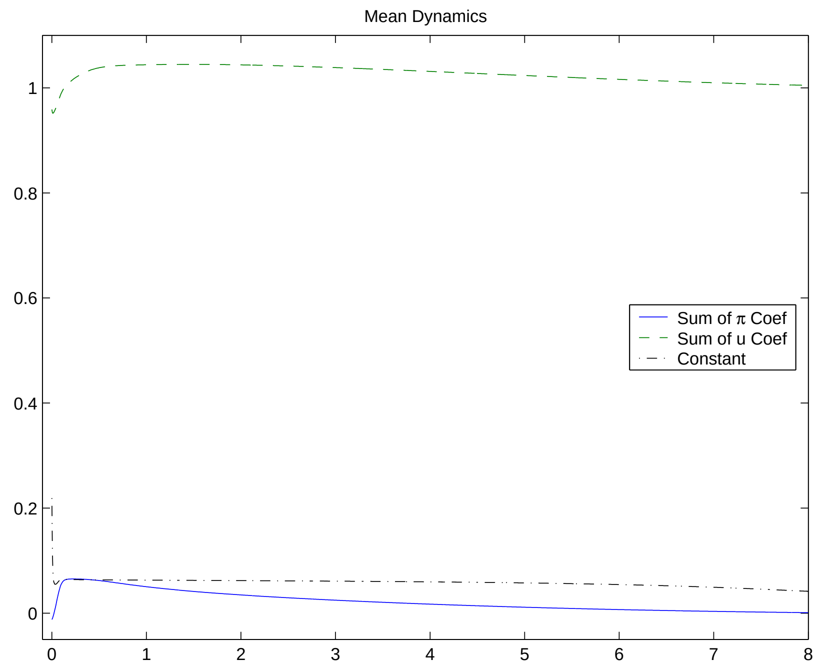

FIGURE 18. Mean dynamics for the baseline estimates, initialized at the start of the sample.

VII. MODEL PROPERTIES

We have shown that our model is capable of tracking the post-WWII inflation data. It is equally important, however, to examine the long run properties of the model to see if the government's adapting beliefs will eventually support near-Ramsey outcomes. We first use the small variation limits of Sargent and Williams (2005) to obtain analytical asymptotic results. Then we discuss the convergence of our baseline model to a limit distribution.

VII.1. **Small Variation Limits.** While it is difficult to obtain explicit convergence results for arbitrary V, for small V's the beliefs drift at a slower rate, allowing us to approximate their evolution with a differential equation. In particular, as in Sargent and Williams (2005), we let $V = \varepsilon^2 \hat{V}$ and study limits as $\varepsilon \to 0$. However, $P_{t|t} \to 0$ as $\varepsilon \to 0$, so we define a scaled matrix $\hat{P}_{t|t} = P_{t|t}/\varepsilon$ that does not vanish. Sargent and Williams show that as $\varepsilon \to 0$, the sequence $\{\alpha_{t|t}, \hat{P}_{t|t}\}$ generated by (9)-(10) converges weakly to the solution of the following

ODEs:

$$\dot{\alpha} = PE\left[\Phi_t(u_t - \Phi_t'\alpha)\right] \tag{17}$$

$$\dot{P} = \sigma^{-2}\hat{V} - PE(\Phi_t\Phi_t')P, \tag{18}$$

where the expectations are calculated for fixed α. As we let the prior belief variance go to zero by shrinking ε, the government's beliefs track the trajectories of these differential equations. We call the ODEs (17)-(18) the *mean dynamics* because they govern the expected evolution of the government's beliefs. If the ODEs have a stable point, then the government's beliefs will converge to it as $\varepsilon \to 0$ and $t \to \infty$. Note from (17) that the limiting beliefs satisfy the key least squares orthogonality condition (5) and hence comprise a self-confirming equilibrium. This orthogonality condition is the key identifying assumption in the government's subjective model, and in the limit it is satisfied when the data are generated by the true DGM.

In Figure 18 we plot trajectories of the mean dynamics for some functions of the parameters describing the government's beliefs starting from the initial conditions at the beginning of the sample. Evidently, the mean dynamics converge to a stable self-confirming equilibrium. The self-confirming equilibrium beliefs are:

$$\alpha_{\text{SCE}} = [\; -0.0008 \quad -0.0000 \quad 0.9725 \quad 0.0000 \quad 0.0165 \quad 0.0688 \;].$$

In the SCE, the government knows the true value of θ_0, the effect of current inflation on unemployment. In the SCE, the government believes in a small tradeoff between inflation and unemployment, and so sets inflation only slightly above the Ramsey level. In particular, the mean inflation rate in the SCE is 2.24% instead of the Ramsey level of 2%.

However, mean dynamics around a self-confirming equilibrium govern the dynamics of our model only for small ε. In practice, for our parameterization ε must be quite small, on the order of 10^{-4}, for the asymptotic approximations to be accurate. Thus, for our baseline estimated V the mean dynamics do not fully characterize the evolution of beliefs. Loosely speaking, for any V and $\varepsilon > 0$, we get convergence to a nontrivial limit distribution of beliefs. Only as $\varepsilon \to 0$ does this limit distribution converge to a self-confirming equilibrium.

VII.2. **Convergence to Near-Ramsey.** What are the long-run implications of the estimated V? The large estimated value of V suggests that one would expect escapes from SCE to be frequent even if the inflation rate at the Nash equilibrium were much higher. To illustrate this point, we change θ_0 from its estimated value of -0.0008 to -1.0 , the value used by Sargent (1999), while keeping all other parameters fixed at the values we estimated. This implies that the Nash inflation rate is around 10%, while the socially optimal Ramsey level remains at 2%. As can be seen from Figure 19, inflation tends to be high, but the large time-variation of the drifting beliefs implied by our estimated V allows the dynamics to escape to low inflation repeatedly, and there is no tendency for inflation to stay for long at the high level. Thus, our V matrix is consistent with repeated escapes in the long run, but they are difficult to detect under our estimates because our estimate of θ_0 implies such a low sacrifice ratio.

To elaborate on this point, Figure 20 shows the inflation dynamics for simulations of 30,000 months starting at different estimated initial conditions: 1960:03 (the beginning of the sample), 1973:12 (the date when both inflation and the perceived trade-off are quite high), and 2003:12 (the end of the sample). Clearly, they all converge to a limiting distribution around the Ramsey outcomes. This convergence occurs from the estimated initial conditions at any date. The fluctuations at the beginning of the simulation reflect the rise and fall of an American inflation process that was temporarily off the SCE equilibrium. As shown in the lower left panel, we are likely to see some high inflation in the near future but such high inflation will be caused purely by exogenous random shocks to inflation, so long as the government continues to see no tradeoff between inflation and unemployment (see the lower right panel of Figure 20). The government's beliefs are volatile for a while but

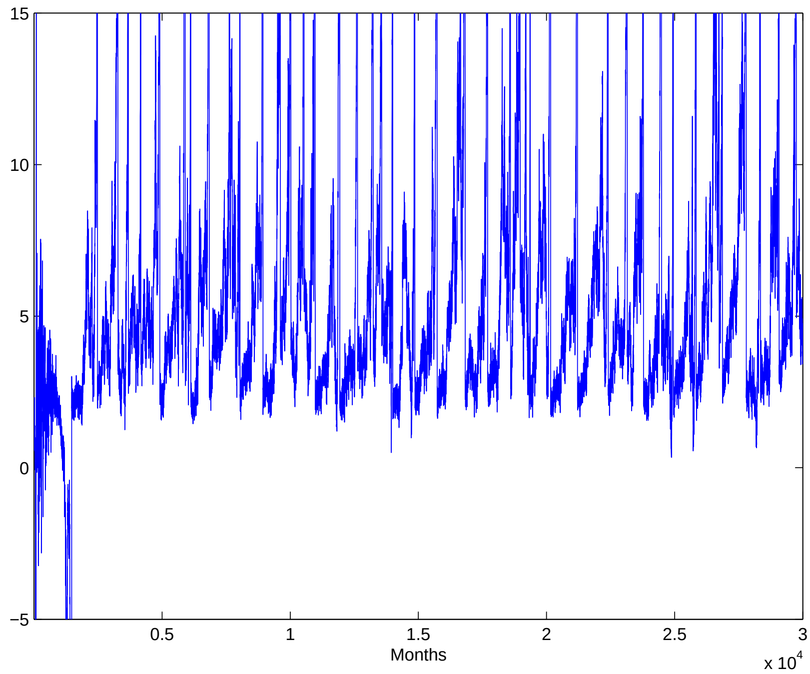

FIGURE 19. Long-run inflation dynamics with $\theta_0 = -1.0$, using the end-of-sample estimate as the initial condition.

eventually the sum of coefficients on inflation converges to near zero.[38] Consequently, the mean dynamics suggest that inflation converges to around 2%. These long run properties foster a view of US monetary history as a process of continual learning before inflation becomes stable around the Ramsey outcomes.

[38] In those volatile periods, the constant coefficient in the government's estimated Phillips curve is often very large (on the order of 100) and the sum of the unemployment coefficients tends to be negative. Thus, these two government Phillips curve parameters fall in the north-west quadrant of the graph discussed in Section VI.1. If the sum of the inflation coefficients is negative, one can see from (16) that the government's dynamic programming problem implies a large increase in inflation to restrain adverse fluctuations in unemployment. Similarly, if this sum is positive, the government tends to generate a large rate of deflation. Such values for the government parameters in our simulations are far outside of the range attained by the historical estimates, as shown in the third row of graphs in Figure 6. When by chance we draw a sequence of shocks that keeps these government Phillips curve parameters within their historical range, convergence to a stable inflation path occurs without large swings of inflation.

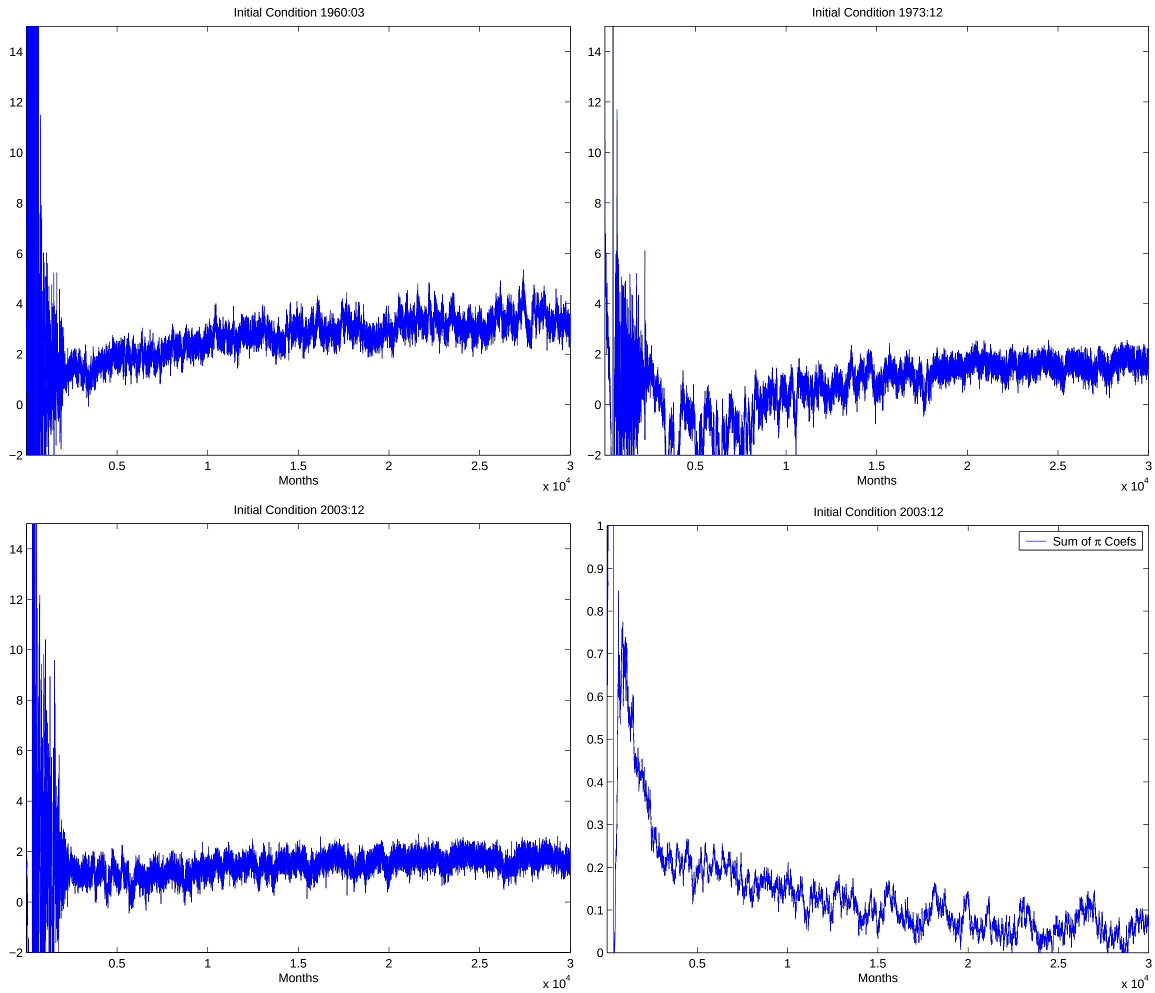

FIGURE 20. Government's inflation choice in long Monte Carlo simulations, using the different estimated initial conditions.

VIII. CONCLUSION

Our estimates attribute the differing inflation outcomes over the post-war period to changes over time in the monetary authority's beliefs. Our empirical results suggest an interpretation that differs from the work we build on. Sargent (1999) and Cho, Williams, and Sargent (2002) suggested that US experience could be explained by convergence to a high Nash inflation level coupled with occasional escapes to a lower Ramsey level. As discussed by

Sargent and Williams (2005), these outcomes also occur in our model when we arbitrarily set parameters of the true Phillips curve to allow a larger gap between the Nash and Ramsey levels of inflation, and when we also impose what, relative to our estimates, is a scaled-down innovation volatility matrix V in the government's belief-drift dynamics (8). However, with our estimates, it appears that oscillations between the Nash and Ramsey levels of inflation, driven alternately by the mean dynamics and then the escape dynamics of the Sargent-Williams (2005) model, were not the main forces that accounted for the inflation process that the monetary authorities in the US chose to administer during the post WWII years. Our estimates of the Nash level of inflation are so near the Ramsey level that it arrests most of the action coming from the mean dynamics and the escape dynamics.[39] Instead, the rise in inflation was driven by the interaction of government beliefs and the shocks impinging on the economy, and the fall in inflation was ultimately due to changes in those beliefs. If the US monetary authorities remember the lessons that prompted Volcker to disinflate in the early 1980s, then maybe Volcker's conquest of US inflation, sustained by Greenspan, will endure.

Appendix A. Data

The two monthly series employed in this paper are:

- Civilian unemployment rate, 16 years and older, seasonally adjusted (source: BLS);
- PCE chain price index (2000=100), seasonally adjusted (source: BEA).

Inflation is measured as an annual rate (12-month ended) of change of the PCE price index. The estimation sample (including lags) is from January 1960 to December 2003.

[39]Furthermore, if we arbitrarily set the parameters of the true expectational Phillips curve to create a big gap between the SCE-Nash and the Ramsey inflations, but retain estimated innovation covariance V, escapes from a SCE again occur frequently enough to vitiate any pattern of recurrent oscillations between the SCE and Ramsey levels of inflation.

APPENDIX B. PRIOR SETTINGS

Our estimation results are quite similar to the maximum likelihood estimates. But the prior is essential for obtaining finite-sample inferences because the government belief parameter matrix V may not have a proper density function when there is no prior. The prior for θ is mostly based on economic theory. For example, the mass prior probability of θ_0 is in the negative region.

The prior mean for θ is set to

$$\begin{bmatrix} 0.12 \\ -0.20 \\ -0.16 \\ 0.98 \end{bmatrix},$$

which implies that the natural rate of unemployment is 6.0 with somewhat persistent unemployment. The prior mean of θ_1 is only slightly less than that of θ_0 in absolute value ($.16 < .20$), implying the low serial correlation of structural disturbances in Sargent's version of the Phillips curve (pp.70-71, Sargent 1999). The prior variance for θ is

$$\lambda_1 \begin{bmatrix} 0.06^2 \\ 0.10^2 \\ 0.08^2 \\ 0.01^2 \end{bmatrix},$$

where λ_1 controls the tightness of the prior variance. With $\lambda_1 = 1$, the prior standard deviation allows large variation but at the same time gives little probability to negative values of v^*, or positive values of θ_0 and θ_1, or the value of τ_1 being greater than 1 (an explosive root).

For the prior of ζ_1 and ζ_2, we set $\bar{\alpha} = 4$ and $\bar{\beta} = 12.5\lambda_2$. By setting $\lambda_2 = 1$, the prior mean for ζ_i becomes 50 and the prior variance becomes 25^2, implying a quite loose prior for ζ_i.

The prior mean for C_P and C_V is 0. The prior variance is $5^2\lambda_3$ for the diagonals of C_P and C_V and 2.5^2 for the off-diagonal elements. The tightness control hyperparameter is set at 0.5.

In this paper, we have checked the robustness of our estimated results by varying the values of the tightness control parameters λ_1, λ_2, and λ_3.

APPENDIX C. CONDITIONAL POSTERIOR DISTRIBUTIONS

Because x_{t-1} does not depend on θ, ζ_1, and ζ_2, it can be seen from (11)–(14) that the posterior distribution of θ conditional on all other parameters is Gaussian and that the posterior distribution of ζ_1 and ζ_2 is of Gamma. Algebra leads to the following propositions.

Proposition 1.

$$p(\theta \mid \mathscr{I}_T, \zeta_1, \zeta_2, \varphi) = \text{Normal}\,(\tilde{\theta}, \tilde{\Sigma}_\theta), \tag{C1}$$

where

$$\tilde{\Sigma}_\theta^{-1} = \zeta_1 \sum_{t=1}^{T} (y_t y_t') + \bar{\Sigma}_\theta^{-1},$$

$$\tilde{\theta} = \tilde{\Sigma}_\theta \left(\zeta_1 \sum_{t=1}^{T} (u_t y_t) + \bar{\Sigma}_\theta^{-1} \bar{\theta} \right),$$

$$y_t = \begin{bmatrix} 1 & z_{2t} & z_{2t-1} & u_{t-1} \end{bmatrix}',$$

$$z_{1t} = u_t - u^* - \theta_0(\pi_t - E_{t-1}\pi_t) - \theta_1(\pi_{t-1} - E_{t-2}\pi_{t-1}) - \tau_1(u_{t-1} - u^*),$$

$$z_{2t} = \pi_t - x_{t-1}.$$

Proposition 2.

$$p(\zeta_1, \zeta_2 \mid \mathscr{I}_T, \theta, \varphi) = \text{Gamma}\,(\tilde{\alpha}_{\zeta_1}, \tilde{\beta}_{\zeta_1})\,\text{Gamma}\,(\tilde{\alpha}_{\zeta_2}, \tilde{\beta}_{\zeta_2}), \tag{C2}$$

where

$$\tilde{\alpha}_{\zeta_1} = \tilde{\alpha}_{\zeta_2} = \frac{T}{2} + \bar{\alpha},$$

$$\tilde{\beta}_{\zeta_i} = \frac{1}{0.5\sum_{t=1}^{T} z_{it}^2 + \bar{\beta}^{-1}}, \quad \forall i \in \{1,2\}.$$

The government's optimization problem renders the conditional posterior pdf

$$p(\varphi \mid \mathscr{I}_T, \theta, \zeta_1, \zeta_2)$$

one of nonstandard form. To draw from this distribution, therefore, we use the following Metropolis algorithm.

Metropolis Algorithm. We employ four steps to simulate φ from its conditional posterior distribution.

(1) Given the value φ^{last}, compute the proposal draw

$$\varphi^{\text{prop}} = \varphi^{\text{last}} + \xi,$$

where ξ is randomly drawn from the normal distribution with mean zero and covariance $\mathfrak{c}\tilde{\Sigma}_{\varphi}$ specified in (D1). The scale factor $\mathfrak{c}$ will be adjusted to keep the acceptance ratio optimal (around $25\% - 40\%$).

(2) Compute

$$q = \min\left\{\frac{p(\varphi^{\text{prop}}|\mathscr{I}_T, \theta, \zeta_1, \zeta_2)}{p(\varphi^{\text{last}}|\mathscr{I}_T, \theta, \zeta_1, \zeta_2)},\ 1\right\}.$$

(3) Randomly draw ν from the uniform distribution $U(0,1)$.

(4) If $\nu <= q$, accept φ^{prop} as the value of the current draw; otherwise, keep φ^{last} as the value of the current draw.

It follows from Propositions 1 and 2 and the properties of the Metropolis algorithm that a large number of MCMC samples alternately drawn from these conditional posterior distributions will eventually form an empirical distribution of ϕ that emulates the posterior distribution.[40]

Appendix D. Proposal density for the Metropolis algorithm

The key to the Metropolis algorithm for the posterior distribution φ is to obtain the covariance matrix for a normal proposal density. Since x_{t-1} is a function of φ, one can approximate it by a second-order Taylor expansion at the posterior estimate $\hat{\varphi}$. It can be seen from (15) that this approximation leads to the following covariance matrix for φ:

$$\begin{aligned}\tilde{\Sigma}_{\varphi}^{-1} =&(\zeta_1\theta_0^2+\zeta_2)\sum_{t=2}^{T}\frac{\partial x_{t-1}(\hat{\varphi})}{\partial\varphi}\frac{\partial x_{t-1}'(\hat{\varphi})}{\partial\varphi}+\zeta_1\theta_1^2\sum_{t=2}^{T}\frac{\partial x_{t-2}(\hat{\varphi})}{\partial\varphi}\frac{\partial x_{t-2}'(\hat{\varphi})}{\partial\varphi}\\ &+\zeta_1\theta_0\theta_1\sum_{t=2}^{T}\left[\frac{\partial x_{t-1}(\hat{\varphi})}{\partial\varphi}\frac{\partial x_{t-2}'(\hat{\varphi})}{\partial\varphi}+\frac{\partial x_{t-2}(\hat{\varphi})}{\partial\varphi}\frac{\partial x_{t-1}'(\hat{\varphi})}{\partial\varphi}\right]+\bar{\Sigma}_{\varphi}^{-1},\end{aligned} \tag{D1}$$

where $\bar{\Sigma}_{\varphi}$ is the prior covariance matrix for φ.

[40]For each draw of ϕ, ζ is normalized to be equal to ζ_1 before the government's inflation policy is solved. This normalization is consistent with Wald normalization discussed in Hamilton, Waggoner, and Zha (2004).

REFERENCES

BLINDER, A. S. (1998): *Central Banking in Theory and Practice*. MIT Press, Cambridge, Massachusetts.

CHAPPELL, H. W., R. R. MCGREGOR, AND T. VERMILYEA (2005): *Committee Decisions on Monetary Policy: Evidence From Historical Records of the Federal Open Market Committee*. The MIT Press, Cambridge, Massachusetts.

CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.

CHUNG, H. (1990): *Did Policy Makers Really Believe in the Phillips Curve? An Econometric Test*. University of Minnesota, Minneapolis, USA, ph.d. dissertation edn.

COGLEY, T., AND T. J. SARGENT (2005a): "The Conquest of U.S. Inflation: Learning and Robutsness to Model Uncertainty," *Review of Economic Dynamics*, 8, 528–563.

——— (2005b): "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.," *Review of Economic Dynamics*, 8, 262–302.

GEWEKE, J. (1999): "Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication," *Econometric Reviews*, 18(1), 1–73.

HAMILTON, J. D., D. F. WAGGONER, AND T. ZHA (2004): "Normalization in Econometrics," Federal Reserve Bank of Atlanta Working Paper 2004-13.

IRELAND, P. (1999): "Does the Time-Consistency Problem Explain the Behavior of Inflation in the United States?," *Journal of Monetary Economics*, 44(2), 279–292.

IRELAND, P. N. (2005): "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," Manuscript, Boston College.

KING, R. G., AND M. WATSON (1994): "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," *Carnegie-Rochester Conference Series on Public Policy*, 41, 157–219.

KYDLAND, F. E., AND E. C. PRESCOTT (1977): "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85(3), 473–492.

PARKIN, M. (1993): "Inflation in North America," in *Stabilizing the World Economy in the 1990s*, ed. by K. Shigehara. Bank of Japan, Tokyo.

PRIMICERI, G. (2003): "Why Inflation Rose and Fell: Policymakers' Beliefs and US Postwar Stabilization Policy," Northwestern University.

——— (in press): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*.

ROBERT E. LUCAS, J. (1973): "Some International Evidence on Output-Inflation Trade-offs," *American Economic Review*, 63(3), 326–334.

ROMER, C., AND D. ROMER (2002): "The Evolution of Economic Understanding and Postwar Stabilization Policy," in *Rethinking Stabilization Policy*, pp. 11–78. Federal Reserve Bank of Kansas City, Kansas City.

SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton, New Jersey.

SARGENT, T. J., AND N. WILLIAMS (2005): "Impacts of Priors on Convergence and Escapes from Nash Inflation," *Review of Economic Dynamics*, 8(2), 360–391.

SIMS, C. A. (1988): "Projecting Policy Effects with Statistical Models," *Revista de Analisis Economico*, 3, 3–20.

SIMS, C. A., AND T. ZHA (1998): "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39(4), 949–968.

——— (2004): "Were There Regime Switches in US Monetary Policy?," Federal Reserve Bank of Atlanta Working Paper 2004-14.

STAIGER, D., J. H. STOCK, AND M. W. WATSON (1997): "The NAIRU, Unemployment and Monetary Policy," *Journal of Economic Perspectives*, Winter, 11(1), 33–49.

NEW YORK UNIVERSITY AND HOOVER INSTITUTION, PRINCETON UNIVERSITY AND NBER, FEDERAL RESERVE BANK OF ATLANTA